Cross-Border Product Adoption: Individual Imports, Migrant Networks, and Domestic Retailers*

David Argente
Yale University and NBER

Esteban Méndez
Central Bank of Costa Rica

Diana Van Patten
Yale University and NBER

March 19th, 2025

Abstract

Globalization and migration are reshaping international trade by linking consumers to distant markets. This paper studies how demand shocks propagate across and within countries, influencing which foreign products consumers buy and retailers import. Using administrative data on *individuals*' foreign purchases and a new instrument that links individual migrants to networks in their home countries, we first show that product-specific demand shocks propagate through international migrant networks. Second, domestically, we show that a new product exogenously imported by a close neighbor, a coworker, or a friend increases an individual's own likelihood of importing that product, especially for more expensive and visible goods. Third, we show that domestic retailers become more likely to import a product if it is popular among consumers who live within its catchment area. We also show, causally and via a large-scale survey, that retailers' responses derive from learning about the local demand for products not yet available domestically by observing consumers' behavior. This mechanism helps retailers identify "preferred" varieties that align with local tastes and benefits lower-income consumers, who do not import directly. Thus, as direct-to-consumer foreign shopping continues to grow, local networks enable retailers to identify the varieties that best match local preferences, changing international trade dynamics.

Keywords: international trade, direct-to-consumer, consumer imports.

---

*Corresponding author: diana.vanpatten@yale.edu. We thank Jason Abaluck, Fernando Alvarez, Costas Arkolakis, David Atkin, Andy Bernard, Matilde Bombardini, Kirill Borusyak, Lydia Cox, Kevin Donovan, Mayara Felix, Teresa Fort, Santiago Franco, Paul Goldsmith-Pinkham, Amit Khandelwal, Sam Kortum, Theresa Kuchler, Reka Juhász, Jeremy Majerovitz, Isabela Manelici, Mushfiq Mobarak, Monica Mrazova, David Nagy, Ezra Oberfield, Luigi Pistaferri, Steve Redding, Richard Rogerson, Todd Schoellman, Pete Schott, Johannes Stroebel, Daniel Xu, and participants at various seminars and conferences for helpful comments and discussion. The views expressed herein are those of the authors and do not necessarily represent the views of the Central Bank of Costa Rica (BCCR). Results have been reviewed by the BCCR to ensure no confidential information is disclosed.

1 Introduction

In recent decades, developing countries around the world experienced a supermarket revolution that shifted consumption from local bodegas and kirana shops to mass-market retailers (Atkin et al., 2018; Reardon et al., 2003). Today, a new transformation is underway: e-commerce is redefining how individuals source their consumption (Bai et al., 2020). While attention has focused on the global rise of online shopping, it is less well known that in remote, small, or developing countries, e-commerce often takes the form of cross-border shopping. For example, in Latin America, consumer online purchases from overseas retailers account for more than 50% of e-commerce in Bolivia, Costa Rica, Ecuador, the Dominican Republic, El Salvador, Panama, and Paraguay, with an average of 39% in the region (EBANX, 2023).

In these markets, consumers face both higher prices and limited product variety due to small scale and geographical isolation. By directly purchasing from overseas retailers, they can bypass these local constraints and gain from variety (Broda and Weinstein, 2006). However, consumers still face information barriers: will the product quality be as advertised? Will the size fit? Moreover, an overwhelming number of options can create a tyranny of choice, leading to suboptimal purchases. In addition, foreign purchases in developing markets are often hard and costly to return, increasing risk for consumers. Therefore, this is a context where demand shocks—originating from consumers' exposure to products—can find fertile ground to propagate and where demand externalities can be key. Similarly, for local retailers, searching among the wide array of options to source products for local markets is costly and risky (Bai et al. 2020; Juhász and Steinwender 2018; Startz 2016), especially when local consumer preferences for these new products are unknown and changing rapidly. Search barriers can be particularly binding for small retailers, who largely determine the assortment available to lower-income households (Faber and Fally, 2022).

This paper studies the role of direct and indirect demand externalities in the process of importing final goods. At the individual level, the role of externalities is direct and intuitive; when an individual buys a foreign product, others around her might become more likely to import it themselves. We explore direct externalities in two ways. First, we study whether product-specific demand shocks propagate across countries via international migrant networks. In other words, if a product is popular in one country, can this trigger imports from countries connected to it via migrants? This force would imply that immigration policy can impact international product diffusion and available varieties. Second, domestically, we analyze how the likelihood of an individual importing a foreign product depends on whether peers in their relevant local network have previously imported the same product, as well as on the product's characteristics.

The third and main part of our analysis focuses on an indirect externality, and leverages both the instrument and the peer-effects results. We aim to understand how retail firms decide what final products to import. In particular, we study whether local retailers can learn about the local demand for products—which are not yet available domestically—by observing which goods are imported by individual consumers and most popular among them. This type of learning—so far unexplored, as the literature has focused on learning along the supply chain (Bai et al. 2020; Startz 2016; Allen 2014, Fernandes and Tang 2014) would help local retailers overcome search barriers and identify varieties that align with local taste, and trigger an impact of foreign demand shocks on local supply responses. As a result of these externalities, the gains from trade would be larger than previously measured, as network effects would amplify policies aiming to stimulate demand via, for example, tariff reductions.

We start by developing a simple conceptual framework to guide our empirical examination—the core of the paper. Our context is the adoption of foreign goods by Costa Rican individuals and firms. Costa Rica is a small open economy where many goods available online are not sold domestically, and international returns are rare.[1] Measuring how consumption of varieties diffuses and whether retailers learn from consumers is non-trivial. First, there are data constraints, as the analysis requires information on consumption and networks. On the consumption front, we leverage novel administrative data on imports by *individual consumers*; for instance, the data would record a purchase at a U.S. online retailer that was then delivered by mail to Costa Rica. Each record includes details like date, price, weight, product code up to the HS-10 level, and country of origin. Using this previously unstudied data, we document new facts on individual imports, including main products imported, origins, and gravity parameters. On the networks front, previous research has mainly iden-

[1] In developing markets, the costs of returning an item are often too high; anecdotally, consumers usually absorb the cost of internationally shipping back the item, the cost of processing the return, and suffer the delays of international shipping. In fact, in the data, as little as 0.01% of individual imports are returned.

tified networks using shared characteristics, such as race or cohort, or has focused on a single network definition.[2] We directly observe networks of relatives, neighbors, coworkers, and friends; and can identify Costa Ricans living abroad and *link them back* to their networks still in Costa Rica. These local networks are computed for the *entire adult population* using information on family trees, employer-employee records, a new measure of friendships developed for this paper (leveraging the universe of local peer-to-peer money transfers), and both local and foreign residences. Crucially, individual-level data on networks and imports can be linked. Finally, to speak to retailers' responses, we use administrative data for all formal retailers, including product-specific imports and catchment areas based on their customers' location. We also conducted a large-scale survey spanning 4% of retailers to complement our analysis and delve deeper into the mechanisms at play.

Second, identifying demand externalities is often challenging (Brock and Durlauf, 2001; Manski, 1993; Moffitt, 2000), as it is hard to distinguish a true network effect from correlated shocks and common characteristics. To overcome these challenges, we propose a new instrument which leverages several aspects of our context and our data, and that is based on the following idea:

> Suppose individuals L and N live in Costa Rica, and L has a sister living in Los Angeles while N has a sister living in New York City. *If product i becomes more popular in LA as compared with NYC in period t, then L is more likely than N to import product i in period $t+1$.*

The spirit behind this instrument is how, anecdotally but also intuitively, information on products is transmitted to developing countries after relatives migrate to developed countries, where more products are available. This instrument exploits that (i) we can identify Costa Ricans living abroad and where they reside in the U.S. (1% of the Costa Rican population), (ii) we can link these immigrants to their family network still living in Costa Rica (5% of the population), (iii) we collected regional data from several sources to follow product-specific dynamics across the U.S., (iv) consumer trends in the U.S. do not respond to local conditions in Costa Rica, and (v) we can track Costa Rican product-specific foreign purchases at a daily frequency. This strategy also has the large advantage that the analysis can be run using product-level

[2] E.g., Agarwal et al. (2018); Bandiera et al. (2009); Charles et al. (2009); Conley and Udry (2010); Kuhn et al. (2016); Maurer and Meier (2008).

individual consumption, as opposed to total consumption, which aids in separating the consumption network effect from income shocks, and will allow us to explore heterogeneous effects. The instrument strongly predicts product-specific imports by individuals with close relatives in the U.S., despite our focus on *residualized* expenditures, which we show are uncorrelated over time and capture the entry and exit of new product brands and varieties in the U.S. Therefore, the first insight from the instrument, and the paper's first contribution, is to show how demand shocks propagate across international migrant networks. Among those with relatives abroad, a one standard deviation increase in exposure to a product leads to an 11-15% higher probability of importing it within a quarter.[3]

We then use the individual-product-time variation from the instrument to examine if, after a Costa Rican with a relative abroad increases her exposure to a product, others in her network (neighbors, coworkers, or friends), but who do *not* have relatives living in the U.S., become more likely to import *the same* product. The 2SLS finds that a 10 percentage point (pp) increase in the share of network members with a relative abroad leads to a 3.5 to 5 pp higher probability of importing this product for individuals without relatives abroad.[4] To grasp the aggregate effect of the direct externalities, suppose there is a $100 increase in U.S. per capita spending on a product. Then, total Costa Rican consumer imports—per individual with a direct connection to the U.S.—would increase by 40 cents.[5] This diffusion channel also has implications for inequality, as families with relatives abroad have, on average, 12% higher incomes than those without such connections. Thus, the demand externality tends to enable lower-income Costa Rican families to benefit indirectly from migration.

We find significant heterogeneity in the way demand shocks diffuse through local networks across products. Given exposure, some products exhibit a strong diffusion, i.e., many people import them, while for others diffusion is weak. Consistent with information being relevant, products diffuse more strongly if they are expensive va-

[3]A one standard deviation is equivalent to increasing per capita spending by $1.5 in the U.S. region where the migrant is located.

[4]As a robustness exercise, we push our data further and construct a distance-3 nodes instrument which extends the methodology of De Giorgi et al. (2019); results hold and remain statistically equal to those of our baseline approach.

[5]The multiplier is in per capita terms based on those with a direct U.S. connection, as they are the bridge through which demand propagates; if the countries were more connected via migrants, the effect would be larger, which is why a per capita estimate is useful.

rieties, visible goods, or belong to more dynamic categories. Suggestively, goods also diffuse more if their initial importer is well-connected and has high centrality.

Finally, our third contribution documents a new channel by which retailers learn about the local demand for foreign products. The limited evidence on the determinants of retailers' sourcing choices—and on whether they learn from their consumers—largely stems from measurement challenges. Even with ideal data, identification remains non-trivial. Moreover, even if consumer learning were causally identified, the mechanisms driving retailers' responses would still be an open question. To make progress, we again leverage our instrument to test whether retail firms with exogenous exposure to a foreign product via their customers are more likely to import that product. In doing so, an added challenge is defining a retailer's catchment area, as retailers may serve multiple neighborhoods. We address this by estimating retailer-specific catchment areas, which we call *retailer gravity zones,* using sales receipt data with customer details (available for roughly two-thirds of retailers) and by proposing a method to approximate the catchment areas for all remaining firms, a strategy that can be replicated in settings lacking customer location data.

With the instrument and gravity zones, we show that retailers respond to the exogenous exposure of their customers to a product. A one standard deviation increase in the share of individuals with relatives abroad who import a product raises the likelihood that nearby retailers import the same product by 9%. This supply response is large and mainly driven by small retailers, who face higher search costs, may benefit more from consumer insights, and might be more receptive to local consumers' needs. Moreover, households who import directly have higher income than those who do not; as retailers introduce varieties of foreign products, the indirect externality expands products access for lower-income households (Faber and Fally, 2022). Thus, while the initial benefits of migration and direct imports are concentrated among the wealthier, diffusion via retail ultimately contributes to variety gains across income groups.

We also explore heterogeneous supply responses based on the product's popularity among individuals. The key is that the peer effects-based analysis conducted earlier allows us to understand the strength of diffusion by product to examine if, conditional on being exposed to a product, retailers become more likely to begin importing it if it featured strong diffusion and *less likely to import it* if the product was not popular among locals. We find that when individuals display *strong* local demand for a foreign good—evidenced by robust diffusion—local retailers respond by importing and selling that product domestically. Conversely, if a product experiences *weak* diffusion after its initial import, indicating low local demand, domestic retailers become *less likely* to import it.[6] Thus, retailers are responsive to the revealed preferences among individuals, in line with learning about local demand for particular foreign goods. Reassuringly, and aligned with the notion that retailers now serve the local market, individual imports of a product decline once retailers begin selling it domestically.

We then design a survey to validate the channels uncovered in the administrative data. Our large-scale survey spans 700 retail firms—approximately 4% of retailers in the country. First, over 80% of respondents receive feedback from their customers on what products to stock. Second, in line with our finding that retailers respond to individual imports, 60% of retailers noted that observing customers importing new products would make them more likely to start importing and selling them domestically. Third, the survey confirms that small retailers rely more heavily on direct customer imports to gauge the local demand for a potential new product. Fourth, how do retailers gauge demand? Around half of those who are responsive to individual imports gain insights into which imported products interest their customers when they physically visit the store and ask about the availability of these goods. Taken together, our findings point to an indirect externality such that retailers—particularly small ones—learn about the local demand for foreign goods by observing the degree of interest for different imported goods across their customers.

Finally, we combine the estimated effects to grasp their overall impact: individual imports, both due to migrant connections and domestic peer effects, and imports of local retailers. Suppose there is a \$100 increase in U.S. per capita spending on a product. In this case, we find total Costa Rican imports of this product—per individual with a relative abroad—would increase in \$5, given the degree of interconnectedness across countries and the strength of the demand propagation.[7] This sizable effect can

[6] Other test in support of firms learning about local demand derives from the notion that employees can be exposed to products where they live and pass information to employers. However, if they live outside the retailer's gravity zone where preferences might differ, their insights should be less informative about latent demand. Indeed, we find that retailers are unresponsive to the exposure of employees who live outside their catchment area, underscoring the importance of local demand knowledge. This finding aligns with survey results, where retailers are five times more likely to gather insights from employees living close-by.

[7] Note that this multiplier is again calculated in per capita terms, where the denominator

be decomposed into additional imports due to the direct externality and individuals' responses—8% of the effect—and the indirect externality and imports of domestic retailers—the remaining 92%. Such magnitudes underscore why accounting for the newly documented supply-side indirect effect when estimating the full response is key.

Related Literature The paper contributes to the literature on trade and information frictions by providing the first direct evidence on how retailers learn from final consumers to inform their sourcing choices and overcome search barriers. This literature has studied price search barriers (Allen, 2014; Steinwender, 2018), search frictions along the supply chain (Chaney 2014; Bai et al. 2020), and product search (Startz 2016; Juhász and Steinwender 2018). Our work closely relates to Startz (2016), who provides insights into how Nigerian sellers overcome search frictions by traveling to find products, and Juhász and Steinwender (2018), who show information technology improvements are valuable for conveying product characteristics. Our work is also closely linked to Bai et al. (2020), who show information barriers in e-commerce are relevant for sellers serving customers abroad. We complement this work by documenting a new channel through which retailers can overcome search barriers and presenting evidence on the mechanism with self-collected survey data. The analysis also estimates retailer catchment areas using novel data on customer residences. Relatedly, Batch et al. (2024) use credit card data to partition the U.S. into consumer zones. Instead, our gravity zones are retailer-specific, and customer location data is available alongside employer-employee data, which allows us to propose an approximation method yielding estimates with a correlation of over 0.98 with those based on customer locations, and which can be used in other contexts.

More broadly, this paper contributes to the literature on international trade and consumption, which has examined how the availability of retail outlets and online platforms impact consumer welfare (Atkin et al., 2018; Couture et al., 2021) and how small firms shape the assortment available to low-income consumers (Faber and Fally, 2022). We focus on direct consumer imports; a new avenue by trade externalities can impact the varieties available locally and which can greatly shape consumer gains (Broda and Weinstein, 2006). To the best of our knowledge, this is the first paper

in Costa Rica are individuals with a relative abroad, as they act as the bridge through which demand propagates. A more connected set of countries would face an overall stronger propagation, which is why a per capita estimate is informative.

to leverage customs data on individual imports. With the direct-to-consumer market rapidly expanding, and expected to accelerate further with increased global internet penetration and improved transport and logistics, this topic is fertile ground for future research, with contemporary work already studying the increasing value of tariff exemptions on individual imports in the U.S. (Fajgelbaum and Khandelwal, 2024).

Moreover, we study direct externalities as a building block in our retailer analysis. Our peer effects study uses a battery of network definitions for the full population—including new measures of close neighbors (via voting records) and friends (via money transfers)—allowing us to assess economy-wide impacts across multiple product categories. The product variation is key in our identification strategy and allows for heterogeneous effects across characteristics, like visibility and value. These elements add to our knowledge, as previous work has yet to fully examine heterogeneous and economy-wide impacts. De Giorgi et al. (2010) make progress with measures based on household income, but observe only total consumption. Bailey et al. (2022) also make progress by using Facebook data to define networks, but rely on a single product. Unlike previous work, we also propose an interaction between direct and indirect externalities, which leads to retailers learning from consumers' experimentation.

Finally, our work relates to studies on the transmission of ideas via migrant networks and the potential benefits for migrants' countries of origin.[8] We propose a method to link migrants to their foreign city and home networks, and show how migrants can facilitate the diffusion of new products to their origin countries; a result that speaks to the effects of migration policy on local variety growth.

The rest of the paper is organized as follows. Section 2 describes the data used in our analysis. Section 3 presents stylized facts on individual imports. We describe our estimation framework and results on direct externalities in Section 4. Section 5 describes our results for retail firms, and Section 6 concludes.

2 Data

We now describe the battery of administrative datasets used in our analysis. Notably, while the data is anonymized, variables *across* datasets can be linked via unique (pseudonymous) identifiers at the individual level.

[8] E.g., Acosta et al. 2008, Beine et al. 2008 and Agarwal et al. 2018. A related paper that leverages homescan panel data along with a structural approach to study how consumption is affected in destination countries is McCully et al. (2024).

Customs Data We leverage customs records from 2014 to 2019. Each import includes up to a 10-digit HS code, along with information on the amount transacted, the quantity traded, arrival date, and the country of origin.[9] As in other countries, customs records are available for firm-level imports. In addition, if an *individual* imports a good (for instance, if she bought an item from an online retailer in the U.S. which was then shipped to Costa Rica), this transaction is also recorded.[10] The median value of these individual imports is $30. To the best of our knowledge, this is the first paper to leverage this type of customs records despite the fact that, far from an unusual practice, cross-border shopping has become prevalent.

Networks Data Within Costa Rica This paper undertakes an effort to combine several reference groups and paint a relatively complete picture of the network of each Costa Rican individual. In particular, we define networks in three different ways.

Networks of neighbors: First, we assume that an individual's network consists of those who live in close proximity. Networks of neighbors are constructed from official records maintained by the National Registry. While records include district of residence, with 488 districts in total, they also detail the voting center which is closest to each citizen's residence *for each adult citizen*. With 2,028 voting centers in total, the median number of adults assigned to each voting center is 586.[11] We propose the latter voting center information as a novel way to get a precise measure of close-by neighbors.

Networks of coworkers: Second, we assume that the relevant network is composed of coworkers.[12] Matched employer-employee data was obtained from the Registry of Economic Variables of the Central Bank, which tracks the universe of formal employ-

[9]For some categories, an HS-10 classification does not exist, so the HS-8 or HS-6 code is the narrowest classification. We use the most disaggregated category available. While HS codes are not barcodes, this can be seen as an advantage in our setting; a person might learn about a new type of flask bottle from a peer, but order a blue one instead of a green one, which would typically be in the same HS code but would not have the same barcode.

[10]While individual imports could potentially also include imports from informal sellers, it will become clear that this would only make our estimates a lower bound: if a person imports a product and informally sells it domestically, then the incentives for others to import it decrease, which would attenuate our coefficient of interest. Moreover, nearly all individuals import each product only once within our sample period (2015-2019), which suggests these are not informal sellers who use this method to stock, and dropping those who import more regularly does not meaningfully change our results.

[11]For details on the distribution of voting centers, see Méndez and Van Patten (2022).

[12]This is in line with De Giorgi et al. (2019), who identify coworkers as a good reference group given the large share of the day spent with them, among other reasons.

ment between 2015 and 2019. These data allows us to recover networks of coworkers which change at a monthly frequency, as people change their employers.

Networks of "friends": Third, we create a novel measure of social networks which connects pairs of individuals who have sent money to each other *bilaterally.*[13] We use data on comprehensive transactions on Sinpe Móvil, an application that since 2015 allows Costa Ricans to make peer-to-peer money transfers via their mobile phones (Alvarez et al., 2023). It has been adopted by nearly 70% of all adults and processes the equivalent of 22% of GDP in transactions each year. We construct a time-invariant measure, as follows: we start at the *end* of the sample period and retrospectively ask: which pairs of peers have made transfers to each other bilaterally? These pairs are considered friends, which has the advantage of eliminating transfers to, for instance, a nanny or a housekeeper. This method allows us to proxy for networks of friends which are usually infeasible to recover; more details are available in Appendix E.

To the best of our knowledge, the breadth of these networks spans more ground than any previous work, enabling us to compare the impact of demand externalities across different networks. Table C.1 presents summary statistics for each network. Networks of neighbors are fewer in number but larger in size, while friend networks are the most numerous and have the lowest median number of members.

Retailer Location and Gravity Zone We leverage data on corporate income tax returns spanning the universe of formal firms in the country. The data span 2015 to 2019, and includes typical balance sheet variables along with details on each firm's sector and location. Section 5 and Appendix I.1 discuss how we construct retailer-specific catchment areas leveraging details on their customers' location, available for a majority of retailers from electronic receipts data.

The instrumental variables strategy proposed in this paper requires three additional data sources, which are described below.

Family Networks We construct *nationwide* family networks based on information from Costa Rica's National Registry. This novel data includes official information to build each person's family tree based on existing records and without relying on

[13]For instance, if user A has only sent money to user B, we would not record this relationship as a friendship. If, however, both A and B have sent money to each other, then their relationship is classified as a friendship.

name-matching. The data is dynamic and at a monthly frequency.

Networks of International Migrants In Costa Rica, voting is mandatory, and it is among the countries with the highest number of immigrants residing abroad who register to vote at their corresponding Costa Rican consulate (approximately 51%).[14] This registration results in a record of migrants' foreign residences maintained by the Supreme Electoral Tribunal. The data maps citizens residing in foreign land to the consulate closest to their residence from 2014 to 2022. Large countries, such as the U.S., have multiple consulates, typically in cities with a high mass of Costa Ricans.[15] While this information is available in other countries, to the best of our knowledge this is the first paper to leverage it to recover international migrant networks.

Consumer Dynamics in the U.S. We obtain U.S. consumer trends by product from two alternative sources, which complement each other. First, we rely on the Consumer Expenditure Survey (CEX), which includes quarterly data by Metropolitan Statistical Area (MSA) for 700 categories of products between 2015 and 2022.[16] While the MSAs for which estimates are produced do not span the entirety of the U.S.'s territory, they do include every city where a Costa Rican consulate is located, which correspond with the main cities where Costa Ricans reside abroad. In fact, according to the American Community Survey, over 82% of Costa Ricans living in the U.S. reside in one of these cities during our sample period. The CEX data (UCC codes) can be then mapped to HS codes using the concordance developed by Furman et al. (2017); also used by Hottman and Monarch (2020) and Borusyak and Jaravel (2021). The variation from this mapping is mainly at the HS-4 or HS-6 level (see Table D.2), as CEX categories are often more aggregated than customs' HS codes.

While our main results are based on the CEX, we leverage a second source of data on consumer trends by product which aims to complement the CEX, precisely by providing variation for narrower product codes. The logic behind this second source is the following: In the U.S., many tradable products are imported. Thus,

[14]For instance, the equivalent share of migrants residing in the U.S. and registered to vote in their home country in Mexico is 1.5%, and the median in Latin America is 17.6%.

[15]The cities with Costa Rican consulates with ratified voting centers are: Atlanta, Chicago, Houston, Los Angeles, Miami, New York, and Washington D.C. There are also honorary consulates in Minneapolis, Puerto Rico, and Tucson. These consulates partition the U.S., and the area serviced by each consulate is well-defined.

[16]Details on geographic coverage are available in the BLS website (link).

U.S. expenditures on these products by region should co-move with the imports of these products in those areas.[17] Following this idea, we use HS-10 level quarterly imports by customs districts in the U.S. from the Census Bureau, which include over 20 thousand product codes, to obtain variation at the HS-10 level (see Table D.2). Conveniently, while HS-10 codes do not necessarily coincide across different nations, U.S. being Costa Rica's main trading partner, they do align for these two countries.[18]

The U.S. has 47 customs districts; instead of assuming a product is consumed in the same customs district it is imported into, we follow Acosta and Cox (2019) and allow for movements of imports within the U.S. using data from the Department of Transportation's Freight Analysis Framework (FAF), which provides estimates of where imported goods travel once they enter into U.S. borders across 132 FAF zones.[19]

Reassuringly, Appendix C.1 presents evidence—both in levels and in *changes*—of a strong correlation between expenditures in the CEX and the one-quarter lagged value of imports by product code and by region. This lag is intuitive, as it takes time both for goods which cross the U.S. border to arrive to retailers and for them to be consumed by households and show up in the CEX; thus, throughout the paper, we use one-quarter lagged U.S. imports as a proxy of contemporary expenditures on those products. Reassuringly, and in line with this strong correlation, we will show that all our main results are statistically equal regardless of whether we measure U.S. expenditures via the CEX or via the U.S. imports data.

To further validate the CEX, Appendix C.1 also leverages transaction-level data on debit card expenditures by region and by Merchant Category Codes (MCCs) in the U.S., with over 10 million cards between 2017 and 2020.[20] While imperfect, as MCCs tend to be designed for financial transaction tracking, it is reassuring that—just as in the case of customs data—the correlation between CEX and card expenditures by region and product code is strong both in levels and in changes.

[17] Note how it is helpful that we will ultimately rely only on variation in (internationally) tradable products, not on changes in expenditures on non-tradables.

[18] Specifically, we manually check that the definition of each HS-10 code that is imported by individuals in Costa Rica is the same as in the U.S.'s Harmonized Tariff Schedule (HTS).

[19] While FAF zones are of moderate size, it will become clear later that, for our purposes, it is not crucial to pinpoint the precise location where a good was consumed; instead, we are interested in the broad area within the U.S. where consumption took place.

[20] These data come from Facteus, a provider of financial data for business analytics.

3 Stylized Facts on Individual Imports

While the determinants of trade between firms have been largely studied, the same is not true for goods imported directly by consumers. This section documents new patterns that govern the decision to import by individuals, and compares them with those in overall trade. These stylized facts, in turn, will be useful to understand the role of peers in individual imports in the next section of the paper.

Which goods are being imported? We first document which are the types of goods that are most commonly imported by individuals. Table A.1 shows a top-10 ranking which results from collapsing imports from HS-10 to HS-4 categories to be more informative.[21] As shown, the most popular category by far is 6204, which includes several types of women's or girl's clothing items. The top categories also include some types of motor vehicles, bags, men's and boys' clothing, toys, and household items. These top categories imported by individuals are very different from the top codes of final goods imported in general at the country-level, and have a modest overlap with the top codes imported by retailers, as reported in Table A.2.

Which origins? Table 1 reports the top origins of imports by individuals, both according to volume of imports and value of imports. The U.S. is the origin of most goods imported directly by consumers, followed by China and Japan. Thus, while the rest of the paper will focus on imports from the U.S., it is worth highlighting that this already captures a large share of the relevant imports. The U.S. is also the main origin for retailers' imports of final goods (35%), followed by China (14%).

Table 1: Top Origins of Individual Imports

Volume		Value	
(1) Origin	(2) Percentage	(3) Origin	(4) Percentage
United States	61%	United States	49%
China	17%	Japan	14%
Japan	2%	China	9%
Bolivia	2%	South Korea	7%
Mexico	2%	Canada	4%

Notes: The table reports the main origins of imports by individuals, both by volume of imports and imports' value. Data spans 2015-2019.

[21] Without such a collapse, most top 10 imports would belong to the 6204 category. Within this HS-4, the most popular HS-6 categories for women are cotton trousers (620462), trousers of synthetic fibers (620463), and dresses of synthetic fibers (620443).

Gravity and Individual Imports We now investigate if individual imports are governed by gravity. For comparison, Table 2 reports the results of estimating a traditional gravity equation both on individual imports and on total imports. Remarkably, the roles of origin GDP and distance for individual imports in column (1) are both statistically equal to the role of the same variables for total imports in column (2) *and* to the role of these variables in the literature (Head and Mayer, 2014). There are, however, a few differences between typical gravity variables in individual imports and those variables in column (2) and in the broader literature. Namely, contiguity seems to play a much more important role for individual imports, and common language, which tends to lead to more total imports, is insignificant (and has a negative coefficient) in the case of individual imports.

Table 2: Gravity Equation for Individual vs. Total Imports (PPML)

	(1) Individual Imports	(2) Total Imports
Origin GDP	1.358	1.247
	(0.113)***	(0.067)***
Distance	-0.889	-0.901
	(0.237)***	(0.126)***
Contiguity	2.968	0.981
	(0.462)***	(0.424)**
Common Language	-0.381	0.721
	(0.331)	(0.296)**
Colonial Dependency	0.303	-0.137
	(2.710)	(0.277)
Time FE	Yes	Yes
Pseudo R^2	0.853	0.922
Observations	925	925

Notes: The table reports the coefficients resulting from a gravity equation with individual imports as a dependent variable in column (1), and with all imports in column (2). The estimation relies on PPML. GDP considers PPP from the World Development Indicators database. Other variables are obtained from CEPII (Conte et al., 2022). Standard errors are clustered by origin country. Data is collapsed to the annual level and spans 2015-2019.

Frequency of importing We also document that, in stark contrast with how retailers import recurrently, nearly all individuals import each HS-10 category only once, also indicating that imports from informal retailers are not widespread. As shown in Figure A.1, repeated imports of the same product are rare.

4 Direct Externalities in Individual Imports

We propose a simple theoretical framework in Appendix B to guide our empirical analysis.[22] The main goal of the paper is to explore these forces empirically. Regarding direct externalities, we want to understand if (i) demand shocks at the product level can propagate across countries via migrant networks, and if (ii) once an individual imports a product, the probability of importing the same product for people in her *local* network meaningfully changes, while exploring what are the products and types of networks for which the effect is the strongest.

4.1 Empirical Strategy

Our empirical strategy leverages several aspects of our setting and data. In particular, we construct an instrument based on the following idea, which we will formalize in the next section: *Suppose both L and N live in Costa Rica, and L has a sister living in Los Angeles (LA) while N has a sister living in New York City (NYC). If product i becomes more popular in LA as compared with NYC in period t, then L is more likely than N to import product i in period $t+1$.*

Panel A of Figure 1 provides a general summary of the instrument. Following the figure's notation, suppose a family in Costa Rica has a member who migrated to a U.S. city on the West Coast, in blue, and a specific product p' becomes more popular in this West Coast city as compared with other cities in the U.S. Then, relatives of this migrant in Costa Rica become more exposed to this product and more likely to import it than other Costa Ricans with relatives in different U.S. cities.

The spirit behind this instrument is how, anecdotally but also intuitively, information on product dynamics is transmitted to developing countries once relatives migrate to developed countries, where more products are available. The instrument exploits (i) that we can identify Costa Rican citizens who are living abroad along with

[22] We extend insights from Caplin and Leahy (1998) and highlight three features: i) the initial delay in the adoption of foreign products, ii) the adoption of these products in a network after someone imports them, and iii) the subsequent adoption by local retailers. The initial delay in product adoption results from its value being uncertain. This delay is not optimal and results from a demand externality; individuals do not internalize that information is revealed once they import. Once someone imports a variety, agents in her network gain information about its type and decide whether to import it. There is also an indirect externality: revealed information triggers retailers' imports, but only if the expected gains are large; firms only import varieties with enough popularity among consumers.

Figure 1: Instrument Relying on International Family Networks and Exogenous Product Trends

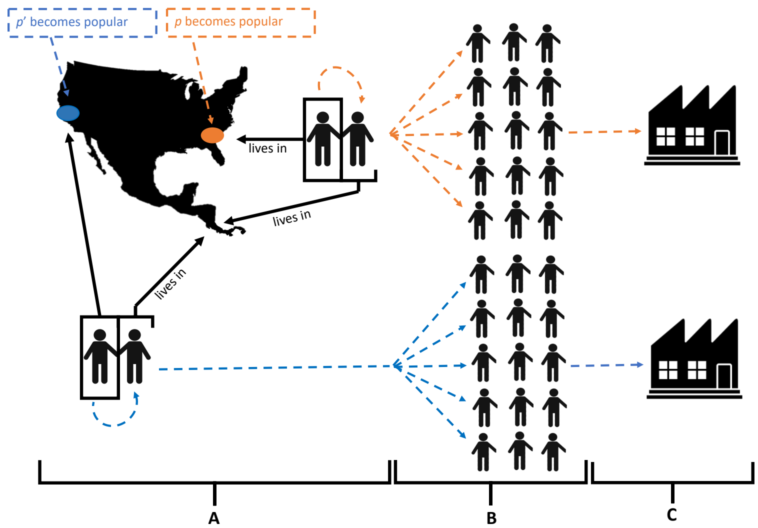

Notes: The figure summarizes the idea behind our main instrument and analysis. The instrument (Panel A) leverages information on the family networks of migrants to different U.S. cities, along with variation on product trends across these cities. Panel B represents our second stage, in which we measure the effect of exogenous exposure to a product on the probability of importing the same product for people who share a network. Panel C represents our study of the supply response after individuals import a product.

the location where they reside in the U.S. (1% of the population), (ii) that we are able to link these migrants to their close relatives who still reside in Costa Rica (5% of the population), (iii) that we have data at the MSA-level and customs district-level to follow product-specific dynamics across the U.S. over time, (iv) that product-specific consumer trends in the U.S. do not respond to local conditions in Costa Rica, and (v) the availability of individuals' product-specific imports at a high frequency. Moreover, the instrument serves a dual purpose: in the first stage, it enables us to test whether product-specific demand shocks propagate across international migrant networks. In the second stage, we leverage it to explore peer effects within Costa Rica (from individuals with relatives in the U.S. to those without such connections).

Demand Shocks in the U.S. We construct a measure of product-specific demand shocks in the U.S., which can vary across time, cities, and products, but that we can purge from business cycles in the U.S., differential level effects, and national product trends. More rigorously, let s denote a U.S. city, p a product, t a quarter, and c a

Costa Rican consulate in the U.S. Consider the following specification:

$$\ln E_{spt} = \alpha + \underbrace{\lambda_{sp}}_{\text{level}} + \underbrace{\mu_{st}}_{\text{local business cycles}} + \underbrace{\phi_{pt}}_{\text{national product trends}} + \widetilde{E}_{spt}, \tag{1}$$

where E_{spt} are expenditures in city s on product p at time t. Let $\widetilde{E}_{spt}$ be residuals of this regression, which would capture the differential product trends across U.S. cities. Note that the fixed effects would prevent $\widetilde{E}_{spt}$ from varying (i) because people in a city are more prone to buy a certain product, for example, those in Chicago buying more winter coats (level effect); (ii) because a particular region had a positive or negative income shock (local cycles); or (iii) because a product became more or less popular (national product trend). In fact, Appendix D.2 draws on more detailed micro-data to show how residuals $\widetilde{E}_{spt}$ are driven by local product dynamics: across products categories, movements in residuals closely follow the entry and exit of product brands in each location. As Costa Rican consulates can span several cities, we aggregate our measure to the consulate-level weighting by population shares in each city:

$$\ln \widetilde{E}_{cpt} = \sum_{s \in c} \left(\frac{CR_s}{CR_c} \right) \ln \widetilde{E}_{spt}, \tag{2}$$

where CR_s/CR_c is the share of Costa Ricans in consulate c who reside in city s.[23] Section 4.7 uses randomization inference in support of these shocks being unconditionally randomly assigned across consulates.[24] Furthermore, as reported in Appendix D.1, we conduct a battery of tests, all of which reject serial correlation in $\widetilde{E}_{cpt}$.

4.2 Demand Shocks Propagate Across Migrant Networks

The first contribution of our paper is to show how product-specific demand shocks propagate across international migrant networks. First, we show this holds at the individual level. Then, we construct an instrument for the first stage of a series of 2SLS analyses to study direct and indirect demand responses.

[23] Costa Rican residents by city are obtained from the American Community Survey. The share is time-invariant as the average Costa Rican residing in the U.S. by 2019 migrated in 1994; thus, movements abroad are rare, and could lead to selection which we prefer to shut down. Therefore, we fix these shares in 2014, one year before our sample period starts.

[24] Moreover, as we show in Section 4.7, a “recentered” version of our exposure measure à la Borusyak and Hull (2023) delivers results which align with those of our baseline exposure.

4.3 Individual-Level Analysis

We first consider the following specification at the individual level, which examines if those in Costa Rica with relatives in the U.S. respond to product-specific demand shocks where their relatives reside:

$$\text{Import}_{ipt}^{US,\, direct} = \alpha_0 + \beta_0 \ln \widetilde{E}_{cp,t-1} + \gamma_{ip'} + \gamma_{it} + \varepsilon_{ipt}, \tag{3}$$

where $\text{Import}_{ipt}^{US,\, direct}$ equals one if individual i *with a relative in U.S. consulate* c imports product p at time t for the first time.[25] Terms $\gamma_{ip'}$ and γ_{it} are individual-product and individual-time fixed effects, respectively—note that a product-time fixed effect would not alter the estimates, as all product×time variation was already removed from the exposure in equation (1). Standard errors are clustered by individual-product.[26]

Results are presented in Table F.1. The regressions, based on over 700 million individual-product-time observations, find that a one standard deviation increase in the exposure to a product leads to a 12% higher probability of importing that product next quarter, compared to other Costa Ricans with relatives in the U.S. The table reports a similar effect using U.S. imports by customs districts to construct exposure.

4.4 Network-Level Analysis and First Stage

We leverage this result as a building block in setting up the first stage of a 2SLS. Importantly, note that any instrument *must* be defined at the network level. To illustrate why, consider an individual in the second stage *without* relatives abroad. Our goal is to determine whether her exposure via others in her network—with relatives abroad—affects her importing behavior. But to do so, we must account for the exposure of *all* individuals in her network with relatives in U.S. cities, aggregating across all cities. Therefore, we now construct a measure of how a Costa Rican network would be exposed to foreign shocks.

[25] A "first-time" import is such if the individual has not imported the HS-10 product since 2005. Throughout the paper, we define a "relative" as including parents, siblings, own children, partner, and partner's parents, siblings, and children.

[26] In $\gamma_{ip'}$, p' controls for either HS-4 codes for the case of the CEX—which is very demanding, as most of the variation is at the HS-4 level (see Table D.2)—and HS-6 codes, for the case of U.S. imports data. To avoid redundancy, and as the network-level result is the one used as a building block in other sections, details on timing, clustering, and the role of each fixed effect will be discussed at length in the next subsection.

Network-Level Exposure We construct measures of exposure at the network-level. The exposure mapping of a Costa Rican in network b to product p via imports follows the popular linear-in-means model:

$$\ln \widetilde{E}_{bpt} = \sum_{c} s_{bc} \ln \widetilde{E}_{cpt}, \tag{4}$$

where s_{bc} is the share of people in network b with a relative living abroad in c.[27]

Network-Level Specification We then consider the following specification:

$$\text{ShareImporters}_{bpt}^{US,\,direct} = \alpha_1 + \beta_1 \ln \widetilde{E}_{bp,t-1} + \gamma_{bp'} + \gamma_{bt} + \varepsilon_{bpt}, \tag{5}$$

where $\text{ShareImporters}_{bpt}^{US,\,direct}$ is the share of people with relatives in the U.S. in network b that import product p at time t for the first time. The terms $\gamma_{bp'}$ and γ_{bt} are network-product and network-time fixed effects, respectively—again, a product-time fixed effect would not alter the estimation as this variation was already taken out in equation (1). The regression is run separately for each type of network, so that $b \in B$ and B is either a neighborhood, a firm, or a friends network; moreover, standard errors are clustered by network-product.[28]

A few remarks are in order. First, this regression only considers imports and exposure of people who reside in Costa Rica *and* have a relative living in the U.S. Second, the left-hand-side variable is conservative as it includes only first-time imports.[29] Third, the battery of fixed-effects strengthens identification. Namely, addresses the "friendship paradox" (Aronow and Samii, 2017), and manages interference in network settings (Borusyak and Hull, 2023), as such interference is product-invariant and inherent to the network. Additionally, γ_{bt} controls for correlated shocks (Manski, 1993),

[27]Note that s_{bc} is fixed across time; Costa Ricans living abroad and all networks are set to 2014, a year prior to the start of our analysis. Furthermore, while most people with relatives in the U.S. have all relatives living in the same consulate, there are a few individuals with relatives in different consulates. For these cases, this regression considers a weighted sum of relatives as the main regressor. Namely, $\sum_c s_{nc} \ln \widetilde{E}_{cp,t-1}$, where s_{nc} denotes n's relatives who reside in consulate c as a share of all her relatives who live in the U.S.

[28]Appendix F.1 explains why, in our particular setting clustering by network-product is sufficient and even on the conservative side.

[29]Note that measurement error on the left-hand-side variable would in general not bias this coefficient. This result holds as long as the exposure's residual is uncorrelated with the measurement error, which in our case is likely to occur.

underscoring how multiple products help resolve challenges unique to single-product settings. On its part, $\gamma_{bp'}$ controls for the relevant product level variation and for network taste and characteristics, for instance, addressing if a rich neighborhood tends to import a product; more generally, it would account for exogenous effects (Manski, 1993). While we saturate the regression with fixed effects, significant variation remains; to visualize it, we compute $\ln \widetilde{E}_{bp,t}$ *netted of fixed effects* in equation (5), and calculate its variance for each network-product pair; Figure 2 shows this variation across networks and products. Finally, the timing of equation (5) is guided by local projection exercises and survey data on typical import durations (Appendix F.6).

Figure 2: Identifying Variation by Network-Product Pair

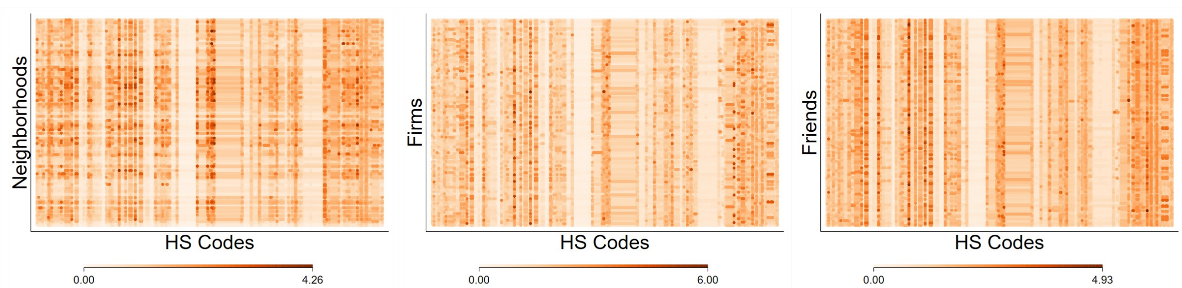

(a) Neighbors (b) Coworkers (c) Friends

Notes: The figure displays differences in the identifying variation across network-product pairs. We compute the term $\ln \widetilde{E}_{bp,t}$ netted of fixed-effects and calculate the variance of this term for each network-product pair. Each panel shows this variance for a network type. Given the large number of networks, for visual purposes panels (b) and (c) are based on a random sample; details in Appendix F.5.

The results of this first-stage are shown in Table 3.[30] Instruments are strong for every network, as reflected by the F-statistics. Moreover, results are remarkably similar across networks of neighbors, coworkers, and friends: A one standard deviation increase in exposure to a product leads to a 11-15% higher share of individuals *with relatives abroad* importing this product next quarter.[31] Table F.4 displays analogous results relying on U.S. imports by customs districts to construct the instrument; reassuringly, they are statistically equal to the baseline results based on CEX data.

[30] Appendix F.5 provides details on the samples of products used in each regression.

[31] A one standard deviation is equivalent to an increase of \$1.5 per capita in the average U.S. consulate. Note that regressions control for network-time fixed effects, thus, for instance, network size would not affect our coefficients.

Table 3: First-Stage Regressions

Dep. variable: Share of importers of product p with a relative in the U.S. and who belong to network b at time t

	%Δ w.r.t. mean import probability		
	Neighbors	Coworkers	Friends
	(1)	(2)	(3)
$\ln \widetilde{E}_{bp,t-1}$	10.676	15.108	10.708
	(1.860)***	(4.630)***	(2.708)***
F-statistic	32.95	10.65	15.63
Observations	289,340,892	300,246,690	260,952,672
Clusters	200,308	237,065	4,568,240
Mean import prob. $[i, bpt]^{US}$	0.001	.0003	0.001
Mean import prob. $[bt]^{US}$	0.150	0.114	0.447
bp, bt, i FE	Yes	Yes	Yes

Notes: The table shows our first stage results. Robust standard errors, adjusted for clustering by network-product, are in parentheses. The independent variables are standardized. We include network×product, network×time, and individual fixed-effects. Percentage mean import probabilities are reported. Appendix F.5 presents details on the sample used in each regression. Data is quarterly and spans 2015-2019.

4.5 Importing Externalities Across Domestic Networks

Relying on the instrument above, we want to understand if people in a network *who are unrelated to migrants in the U.S.* increase their probability to import a particular product after being exposed to it via their peers who do have relatives abroad (Panel B of Figure 1). Thus, from our IV's first stage, we leverage predicted values for the shares of importers of a product *with relatives in the U.S.* as explanatory variable. Our dependent variable would instead depend on the probability of importing a particular product for people in the network without relatives in the U.S., as follows:

$$\text{Import}_{i,bpt} = \alpha_1 + \beta_2 \, \widehat{\text{ShareImporters}}_{bp,t-1}^{US,direct} + \gamma_i + \gamma_{bp} + \gamma_{bt} + \varepsilon_{i,bpt}, \quad (6)$$

where $\text{Import}_{i,bpt}$ equals one if individual i in network b *without relatives in the U.S.* imports product p at time t *for the first time*, and where we again include a battery of fixed-effects so that we only exploit bpt-level variation. We also include an individual fixed effect, γ_i.[32] Just as for the first stage, we consider three network definitions: neighbors, coworkers, and friends, and run independent regressions for each of them. Each network has complementary strengths. Networks of neighbors span all the population and will allow us to study the role of indirect demand externalities in

[32] As there are millions of Costa Ricans *without* relatives abroad and a 2SLS with billions of observations is unfeasible to run, we use a random sample as detailed in Appendix F.5.

triggering a supply-side response. Networks of coworkers might have less correlated spatial shocks, however, they only span the formally employed; 41% of the population. The networks of friends are a novel way of measuring connections beyond observables, but the analysis is limited to those people who have adopted the mobile payment app; 60% of the population. Overall, utilizing all networks paints a better and more robust picture of the role of demand propagation in product adoption.

Table 4: 2SLS: Propagation Within Network

Dep. variable: Prob. importing product p for individual i without relatives in the U.S. and who belongs to network b at time t

	%Δ w.r.t. mean import probability		
	Neighbors (1)	Coworkers (2)	Friends (3)
$\widehat{\text{ShareImporters}}^{US\,direct}_{bp,t-1}$	13.479	18.003	14.679
	(5.218)***	(8.098)**	(5.027)***
F-stat first stage	32.95	10.65	15.63
Observations	289,340,892	300,246,690	260,952,672
Clusters	200,308	237,065	4,568,240
Mean import prob. $[i, bpt]^{US}$	.0002	.0001	.0004
Mean import prob. $[bt]^{US}$	0.044	0.052	0.156
bp, bt, i FE	Yes	Yes	Yes

Notes: The table displays the results of running our 2SLS. Robust standard errors, adjusted for clustering by network-product, are in parentheses. The independent variables are standardized. Regressions control for network×product, network×time, and individual fixed-effects. Percentage mean import probabilities are reported. Appendix F.5 details the sample used in each regression. Data is quarterly and spans 2015-2019.

The baseline results of the 2SLS estimations are shown in Table 4.[33] The magnitudes of the 2SLS coefficients are similar across networks; we find that a one standard deviation increase in share of people with relatives in the U.S. who import a product leads to an increase of between 13% and 18% in the individual probability of importing this specific product within a quarter for those in their network without relatives in the U.S. Put differently, a 10 pp increase in the share of those in a neighborhood with relatives abroad leads to a 4.5 pp higher probability of importing this product for individuals in this neighborhood who do not have relatives abroad, with the corresponding effects for coworkers and friends being 3.5 pp and 4.7 pp, respectively—see Table F.3 for results without normalizations. Reassuringly, results when relying on U.S. imports data to construct our instrument are very similar (and statistically

[33]Appendix F.5 provides details on the samples of products used in each regression.

equal) to the baseline results based on the CEX, as shown in Table F.4.[34]

4.6 Aggregate Effect of the Direct Externalities

Mean import probabilities in each regression are small *by design*: They represent the probability that an *individual* in a given network orders *a particular product from the U.S. for the first time* on a determined *quarter*, as the left-hand-side variable is individual-product-time-specific and "turns on" only with first time U.S. imports.[35] To understand the economic relevance of the estimates, one has to consider the increase in import probabilities across millions of individuals, networks, and products, which we do as a back-of-the-envelope exercise.

Combining results from our first and second stages in equations (5) and (6), and defining $Importers^{US,indirect}$ as Costa Ricans without a relative in the U.S., we can obtain a back-of-the-envelope estimate of how U.S. spending on a product would influence total Costa Rican individual imports of that product. Specifically, we compute the following objects following a change in U.S. spending on a product:

$$A = \underbrace{\Delta ShareImporters^{US,direct}}_{\beta_1 \Delta \widetilde{E}_{bpt}} \times Importers^{US,direct} \times AvgPrice^{Individual} \tag{7}$$

$$B = \underbrace{\Delta Prob(Import_i)}_{\beta_2 \Delta ShareImporters^{US,direct}} \times Importers^{US,indirect} \times AvgPrice^{Individual}, \tag{8}$$

where $AvgPrice^{Individual}$ is the average value of an individual import. A in equation (7) would then speak to the value of new imports of directly exposed individuals, as it is the product of the change in the share that imports, the total importers with direct connections and the average price of an individual import. B in equation (8) captures the value of new imports of those indirectly exposed via their network. Using networks of neighbors as our reference, a comparison of B/A in the above equations along with the estimation from column (1) of Table 4 imply that a new exogenous individual import generated by the directly exposed consumer in her network would

[34] Appendix F.3 discusses the ordinary least squares (OLS) and the reduced form results.

[35] While we have run all these regressions at the network-product-quarter level, delivering much larger import probabilities, individual-level regressions are cleaner, more transparent, and better identified. Results at the network level are consistent with individual-level results and available upon request.

lead to approximately 13 new imports by indirectly exposed people in her network.

As an example, suppose that $\Delta\widetilde{E}_{bpt}$ corresponds to a \$100 increase in U.S. per capita spending on a product. In this case, this calculation implies that Costa Rican consumer imports per individual with a direct connection to the U.S. would increase by 40 cents, given the degree of interconnectedness between the two countries and the strength of the demand propagation across individuals.[36] As will be studied in Section 5.3, while this multiplier is non-negligible, the impact of imports by individuals (i.e., $A+B$) will represent only a modest fraction (about 8%) of the total impact on local demand, as the bulk of the increase will emerge from retailers' responses.

4.7 More Demanding Specifications and Robustness

It is worth spelling the exclusion restriction of our instrument. Our identification strategy requires that the likelihood of buying product p of a Costa Rican—without relatives abroad—in a network connected to a U.S. city via family ties co-moves with *changes* in expenditures on p in this U.S. city only through the relatives' influence.[37] Arguably, our main specification, which is saturated with a battery of fixed effects, takes care of most first-order concerns related to this statement. To complement it, we now conduct a series of robustness exercises with yet more demanding specifications.

Instrument Using Distance-3 Nodes We can push our data to construct an alternative instrument which can rule out lingering alternative hypotheses. The instrument exploits that we have information on both coworkers and spouses, and that spouses who work at different firms can be seen as a bridge between sets of coworkers that are otherwise disjoint; an observation in De Giorgi et al. (2019), and that we extend to include product-level variation and an instrument for the 3-nodes-away exposure. The approach relies on the notion that *if the coworker of the spouse of my*

[36]Note that this multiplier is calculated in per capita terms, where the relevant denominator in Costa Rica is the number of individuals with a relative abroad. We consider this to be the relevant parameter, as these individuals with relatives in the U.S. act as the bridge through which demand propagates. A more connected pair of countries (with more directly exposed individuals) would likely face a total stronger degree of propagation, which is why a per capita estimate is informative.

[37]While, given our fixed effects, this statement only has to hold in changes, Appendix D.3 also presents evidence in support of this statement in levels. We find balanced observables (age, gender, wage) among Costa Rican migrants to different U.S. cities.

coworker has a relative in the U.S. and becomes exogenously more likely to import a product—controlling for common shocks experienced at my firm—this should not influence my probability of importing this exact same product directly, only indirectly though peer effects. Figure 3 presents a diagram to make this notion more clear; it

Figure 3: Diagram of Instrument Using Distance-3 Nodes

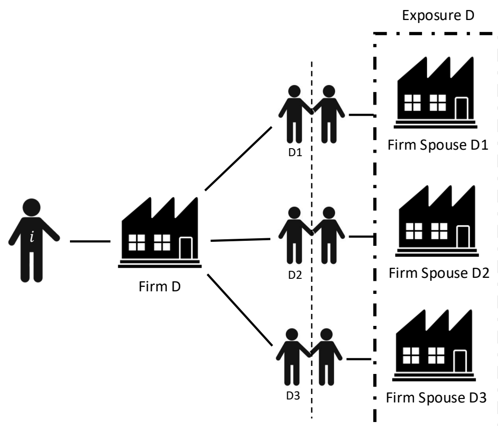

Notes: The figure shows the idea behind our instrument, where the relevant exposure is product-specific, time-varying, and depends on exogenous consumer trends, as described in Section 4.1.

considers an individual i working at firm D. The individual's exposure to a particular product p depends on the exposure to p faced by the spouses of her coworkers (D1, D2, D3) at their firms, which in turn depends on the family ties that employees of those firms have with people residing in different U.S. cities, and on how expenditures of product p evolve in those cities. We then consider the following regression for individual i, which depends on product p at time t:

$$Import_{i,Dpt} = \delta_0 + \theta \widehat{ShareImporters}_{ip,t-1} + \delta_{Dp't} + \delta_i + \varepsilon_{i,Dpt}, \tag{9}$$

where the dependent variable, $Import_{i,Dpt}$, is equal to one if individual i at firm D imports product p at time t for the first time. On the right-hand side of the regression, $\widehat{ShareImporters}_{ip,t-1}$ is instrumented by the mean exposure of firms employing the coworkers' spouses.[38] $\delta_{Dp't}$ are own-firm×product×time fixed-effects; these fixed-effects are key, as they force the identifying variation to come from *differences* between

[38]Note that De Giorgi et al. (2019) rely on characteristics of the firms employing the coworkers' spouses; in our case, this would imply using imports directly. Instead, we opt for a more demanding specification that instruments for those imports.

the shocks at coworkers' spouses firms and shocks directly affecting firm D. Finally, δ_i are individual fixed-effects. This specification is very demanding, however, there is still significant variation left across individual-product pairs, summarized in Figure 4. We also list the product codes which are among the top 10 codes of this subsample but were not in the top 10 codes in the full sample; remarkably, all of these have to do with work-related products. More details are available in Appendix G.1.

Figure 4: Distance-3 Nodes: Identifying Variation by Individual-Product Pair

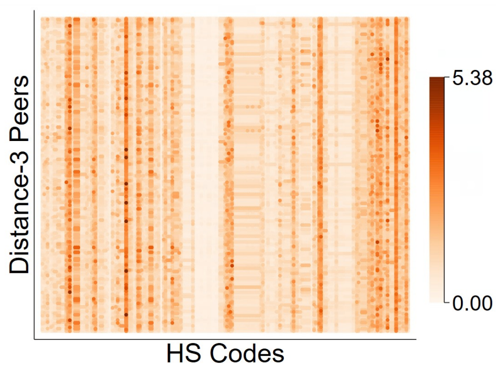

Distinct top codes:
Wooden furniture of a kind used in offices (940330)
Pens and marker parts (960899)
Women's suits of synthetic fibers (610413)

Notes: The figure summarizes the differences in the identifying variation across individual-product pairs. Namely, we compute the term $\widehat{Import}^d_{ip,t}$, netted of the fixed-effects in equation (9), and calculate the variance of this term for each individual-product pair. We also display the product codes which were both among the top 10 codes of this sample and were not included in the top 10 codes in the full sample.

Key advantages of this instrument The distance-3 nodes instrument is immune to several identification concerns. As an example, consider correlated preferences among people in a network and their relatives abroad, i.e., people with relatives in NYC have different product-time demands than those with relatives in Houston, and their friends in Costa Rica show the same differential demand patterns. The concern would be selection into locations, so people go to NYC because the city's preferences are correlated with theirs and because they have Costa Rican friends/colleagues who are similar to New Yorkers.[39] Given this scenario—potentially the worst possible for our baseline instrument—there are two possibilities: (a) there is *not* assortative matching in the marriage market along lines which influence product demands, in which case the instrument based on Figure 3 would deliver a correct estimate; or (b) there is *is* assortative matching in the marriage market along lines which influence product

[39] Note this would have to happen while maintaining balanced observables of migrants across locations (see Appendix D.3).

demands, in which case, the own-firm-product-time fixed effect in equation (9) would co-move with our instrument and would prevent θ from being identified from such assortative matching; again, this approach would deliver a correct estimate.

We find evidence in support of demand externalities, even under this more demanding specification. As shown in column (2) of Table 5, the effect, statistically significant at the 1% level, is an increase of 21% in the probability of the individual importing the specific product within one quarter, with respect to the mean probability of importing. This magnitude is similar to the one documented in column (1), which relies on our baseline IV and the same sub-sample.[40]

Table 5: 2SLS: Individual Imports and Distance-3 Exposure

Dependent variable: $Import_{i,Dpt}$
(Prob. of individual i of importing product p at time t)

	%Δ w.r.t. mean import probability	
	Baseline IV	Distance-3 IV
	(1)	(2)
$\widehat{ShareImporters}_{ip,t-1}$	32.298	21.301
	(10.502)***	(3.617)***
F-stat first stage	34.68	10,787.2
Observations	396,592,974	396,592,974
Clusters	452,235	17,479,088
Mean import prob. $[i, Dpt]^{US}$	.0003	.0003
Mean import prob. $[Dt]^{US}$	0.051	0.051
Dp', Dt, and i FE	Yes	No
$Dp't$ and i FE	No	Yes

Notes: Both estimations are constructed based on the same sample and run at the individual level. Robust standard errors, adjusted for clustering by network-product in column (1) and individual-product in column (2), are in parentheses. The independent variables are standardized. Regressions in columns (1) and (3) control for own-firm×product×time and individual fixed effects, while column (2) controls for individual-product, individual-time, and product-time fixed-effects. Percentage mean import probabilities are reported. Appendix F.5 presents details on the sample used in each regression.

Recentering and Placebo Exposure Measures A key assumption in our setting is that our shocks are randomly assigned. To think about randomization inference, we permute our shocks $\tilde{E}_{cpt}$ *across U.S. consulates, within an HS-4 category and within a time period.* We then run both our first-stage and the reduced form regressions using

[40]The first stage F-statistic in column (2) is large. The reason is that there might not be an endogeneity issue even if we were relying on *imports* of the spouse's coworkers *without* instrumenting with our foreign exposure measure. In fact, De Giorgi et al. (2019) do not instrument and just rely on the distance-3 nodes to be enough to overcome endogeneity issues. Regardless, we believe instrumenting is desirable in our case to make estimates more comparable and to reduce potential measurement error.

the reassigned placebo exposure, and repeat this exercise 100 times—we focus on the reduced form since there will no longer be a first stage for the IV. Figure G.1 plots the distribution of placebo coefficients and depicts the actual coefficient based on the "true" shocks with a vertical red line. The actual coefficients are far in the tails of the placebo distributions. Further, following Borusyak and Hull (2020), we construct a "recentered" version of our exposure measure, by subtracting the expectation of the treatment value under the randomized distribution from our original exposure. As predicted by the theory, this method leads to a slightly larger, but similar, coefficient, as shown in Panel (c) of Figure G.1 for networks of neighbors.

More Demanding Controls Adding certain fixed-effects to our specification can be a powerful tool to rule-out alternative hypotheses. We start by considering a *district*×*product*×*time* fixed-effect, and re-run the analysis defining networks as neighborhoods. Recall that our variation is at the *neighborhood*×*product*×*time* level, so including this fixed-effect limits us to consider variation *within* small areas.[41] Results remain largely unchanged, as shown in column (1) of Table G.1. This is useful, for example, to rule out a story where a seller is targeting an area of the country with advertising about a product.

In a similar spirit, we can add a *network*×*HS-2 product code*×*time* fixed-effect to our analysis. Thus, we would only be exploiting variation *within* relatively narrow product categories. As column (2) of Table G.1 shows, effects are again largely unchanged. Like the exercise using distance-3 nodes, this result speaks against people from a certain network having a preference for a product category, and thus moving to cities where this category is trendy.[42] Also like the analysis with distance-3 nodes, this control would take care of sector-level trends in particular cities.

4.7.1 A Remark on Observed Networks

It is not possible to observe all the connections that each Costa Rican has with the U.S. We cannot observe Costa Ricans who are not registered at a U.S. consulate, but beyond this, people might know U.S. residents and communicate with them, even if they are not relatives. This challenge is pervasive in the networks literature

[41] Each district has four neighborhoods on average.

[42] The logic behind the fixed-effect is that, for example, a person might move to NYC because she likes fashion (HS-2), but is unlikely to move because she likes female trousers made of wool.

(Goldsmith-Pinkham and Imbens, 2013; Manski, 1993). Nonetheless, we believe we have enough evidence to show that this is not a major concern. First, note that as long as the measured shares s_{bc} are proportional to the true shares, this concern should *not* impact our first stage results, so our first finding on the propagation across international migrant networks would remain unchanged. To see why, note that s_{bc} in equation (4) would remain unchanged as long as the overall shares by consulate remain the same; for instance, if people have unobserved connections in the same U.S. cities where they have relatives. Another (potentially less likely) possibility lies in the other extreme of the spectrum: the measured connections in U.S. are located in cities which are orthogonal to the ones of the unobserved ones. In this case, the instrument would suffer from classical measurement error, which would only bias results downwards, but in our case does not lead to a weak instrument. Second, while missed connections could impact the interpretation of our second stage, the distance-3 nodes instrument includes own-*network*×*product*×*time* fixed effects—thus it is immune to this concern—and still delivers similar results to the 2SLS in our main specification, indicating that, if anything, any bias is small.[43] Third, as will become clear, our results on retailers depend on (i) our instrument as constructed in the first stage and (ii) the overall *ranking* of product-specific propagation in the second stage. Thus, any bias that re-scales the 2SLS coefficients but preserves the propagation ranking would be inconsequential for the results on retailers.

4.8 Heterogeneity in Demand Propagation

We now explore the determinants of the strength with which a product propagates within a network after it is imported. To do so, we will create interaction terms *which will then be instrumented*. Therefore, at the bottom of each table including an instrumented interaction term, we also report Sanderson and Windmeijer (2016) conditional first-stage F-statistics for the validity of the instruments. For reference, we also report the Stock–Yogo 10 percent and 15 percent critical values for a perfectly identified model with two endogenous variables (7.03 and 4.58, respectively), which are the appropriate thresholds to reject that the instruments are weak.

[43]Further, recall that our first stage includes network-time fixed-effects, which would prevent biases which are constant across products. Moreover, we can relax this condition further, as results hold controlling for *network*×*HS-2 product code*×*time* fixed-effects, which would prevent biases that are constant *within* product categories.

Dynamic Product Categories We start by comparing dynamic product categories with established ones, as one would expect the information channel to be particularly relevant for categories with more dynamism. We rely on Business Dynamics Statistics (BDS) data, which tracks dynamics on establishments with paid employees. These measures are available for the entire U.S. economy, and by industrial sector, 4-digit NAICS, state, and MSA.[44] Specifically, we use data on the creation of jobs by new establishments and on the entry of new establishments by product category to classify a product as "dynamic" ("established") if its creation of jobs by new establishments and entry of new establishments is above (below) the median within our sample (2015-2019).[45] Table H.1 shows our results, which are consistent regardless of the definition of dynamic products. We document a stronger propagation of products in more dynamic categories, as shown by the positive coefficients in the (instrumented) interaction terms. This result aligns with demand externalities aiding in relaxing information frictions, which might be larger in more dynamic categories.

Centrality of the Importer Products might propagate more if they are initially imported by someone more connected to others. To explore this, we create a measure of degree centrality, which depends on how many friends a person has using our app-based definition of friendship.[46] We then consider how diffusion in a network depends on the average centrality of its members *with relatives abroad*. Results in Table H.2 suggest that the more central the importers in the first stage, the stronger the diffusion across the neighborhood in the second stage. While these results are indicative, note that the interaction term is noisy; this aligns with recent findings from Akbarpour et al. (2023), which document that the choice of optimal seeds can have limited impact on diffusion within a social network.

[44]The BDS is created from the Longitudinal Business Database (LBD), a confidential database used by qualified researchers via secure Federal Statistical Research Data Centers.

[45]The two variables constructed are: (i) entry of establishments, which is the share of new establishments over total establishments in a product category; and (ii) employment gains from new establishments, which equals the share of jobs created by new establishments to total employment in the product category. We then define the variable $Dynamic_p$ used in columns (1)-(3)—definition (i)—and columns (4)-(6)—definition (ii)—of Table H.1.

[46]Degree centrality is one of the simplest centrality measures; a node's degree is a count of its friend connections, and the degree centrality for a node is just its degree. For instance, a node with 4 friends would have a degree centrality of 4.

Goods' Visibility Intuitively, demand shocks for more visible goods should propagate more easily; either because others are more likely to inquire about a visible good or gather information on it, or because conspicuous consumption may amplify the propagation of demand shocks for such goods. In any case, the more visible a product is, the more likely it is to propagate (and vice versa for non-visible goods). To test for this force, we rely on the product-specific visibility index developed by Charles et al. (2009). We then construct an indicator variable equal to one if the good is *below* the median visibility score. As shown in column (2) of Table H.2, there is a dramatic difference in diffusion between goods depending on their visibility; for non-visible goods, the instrumented variable has a 37 p.p. weaker effect on the probability of importing compared to visible goods.

Expensive or "Premium" Products We finally explore if results are heterogeneous between types of goods depending on whether they are *premium* or not. We define premium goods as those whose average price per kilogram is above the median of their HS-4 product category.[47] We focus on networks of neighbors and later also explore if retailers are also more responsive to these more expensive varieties. Table H.3 shows the results of interacting a product-specific premium dummy with our exposure measure. Our findings in column (1) show how the direct externality is about twice as large for premium goods. The latter aligns, for instance, with information spreading more for products which are more expensive (and therefore riskier to import ex-ante).

5 Learning from Consumers: Retailers' Imports

So far, we documented direct demand externalities. First, we showed how demand shocks propagate across international migrant networks. We then documented how after an individual imports a product, others in her local network become more likely to import it. We now explore an indirect externality. We study if, once individuals in a network decide to import a product, there might be useful information about the local demand for this product which becomes available to domestic retail firms. For instance, retailers might be more eager to start importing a product the more locals are willing to acquire it, i.e., the stronger the observed propagation after an individual

[47] Results are robust to different cutoffs as long as they are over the 50th percentile.

import. In contrast, retailers might be more cautious about importing products that show weak demand propagation among consumers (Panel C of Figure 1).

Causally documenting this learning from customers is not simple; for instance, it could arise from common shocks or correlated preferences. This challenge has led to limited evidence on this topic in the firm dynamics literature. To test these forces, we leverage both the first *and* second stage results of the previous section. Namely, we begin by considering the following regression, which leverages the instrument based on U.S. connections:

$$\text{Import}_{f,bpt} = \alpha_3 + \beta_3 \widehat{\text{ShareImporters}}^{US,direct}_{fp,t-2} + \gamma_f + \gamma_{bp} + \gamma_{bt} + \varepsilon_{bpt}, \tag{10}$$

where $\text{Import}_{f,bpt} = 1$ if retail firm f in neighborhood b imports product p at time t *for the first time*, γ_f are retailer fixed effects, and other right-hand-side variables are defined as in equation (6).[48] Ideally, the variable $\widehat{\text{ShareImporters}}^{US,direct}_{fpt}$ would depend on retailer f's catchment area—i.e., its area of influence where its customers are—as retailers' decisions are likely to be influenced by their clients. Therefore, we construct retailer-specific catchment areas, that we refer to as *retailer gravity zones* to then create a weighted-average of the exposure faced by *each retailer's own customers.* To do so, we use information on the residence of each retailer's *customers*, which is available for a majority of retailers from data on electronic vouchers. Moreover, we propose a method to approximate these gravity zones for *all* retailers based on employees' residences, which can be applied in other contexts where the customer-specific location data is unavailable and which delivers exposure measures correlated almost perfectly (0.98) with those leveraging customers' location. Appendix I.1 includes details on the gravity zones and exposure construction.

Then, to study if retailers' responses depend on the likelihood of a product to propagate, we classify products depending on *individuals' response after an exogenous import.* We first start with an exploratory exercise in which we run our second stage regression (equation (6)) but with *product-specific* exposure measures, which allows us to recover one coefficient β_p per product and rank products according to

[48] Our panel on imports begins in 2005 while our regression sample runs from 2015 to 2019. We consider an import as a "first-time" import if the retailer has not imported the HS-10 product since 2005. Relaxing this strict constraint to be the first import within the 2014-2019 period does not substantially change the results.

the magnitude of their propagation.[49] We construct an indicator variable that corresponds with the strength of product propagation; $LowProp_p = 1$ if β_p is in the bottom 10th percentile—a threshold chosen so that the dummy focuses on products featuring a negative effect on demand (i.e., negative β_p, as shown in Figure I.1). Then, we construct an interaction term with this indicator and the exposure measure. Now, while the interaction with $LowProp_p$ might be informative, it is based on an endogenous object. Thus, we complement the analysis with an exogenous proxy of $LowProp_p$. Namely, as shown in Section 4.8, less visible products propagate less across individuals. Therefore, we also consider an interaction with the indicator $LowVisibility_p$ in the exploration of the mechanism behind the retailers' responses.

5.1 Results on Retailers and Mechanism

Table 6 has the results of the analysis. First, in column (1), we find that retailers respond to an increase in exposure to a product. A one standard deviation (10 pp) increase in the (instrumented) share of individuals with relatives abroad who import a product raises the probability that a local retailer imports the same product for the first time by 9 pp. Thus, we document that retailers respond to the observed local demand for foreign goods by importing.

We now delve deeper into the mechanism behind this result. Column (2) of Table 6 is informative about this mechanism. This column shows that the effect of individual imports is driven by high propagation goods, which are the ones that tend to trigger a supply response from retailers in the form of a higher likelihood of importing specific products.[50] Instead, if a product exhibits low propagation, retailers *becomes less likely* to import this product. This asymmetric response points to retailers learning about the *level of the local demand* for those products, as opposed to just a product-discovery story. For example, if a foreign product exhibits high propagation, people in the firm's catchment area may go to local stores and inquire if they have the product in stock. This increased interest might lead firms to order the product in the future,

[49]This ranking is only done for products for which β_p is significant at the 10% level. The distribution of β_p is reported in Figure I.1. For example, the highest β_ps correspond with subtypes of telephone sets and varieties of knitted women jackets; while the lowest β_ps correspond with blank compact discs, and distinct types of ACs and wooden seats.

[50]Table I.5 shows that this result remains statistically equal when we allow retailers to import from any origin country, as opposed to considering only imports from the U.S.

Table 6: Supply Response from Retailers

Dependent variable: Prob. of retailer f importing product p at time t

	%Δ w.r.t. mean import probability				
	All Retailers			Small Retailers	Large Retailers
	(1)	(2)	(3)	(4)	(5)
$\widehat{\text{Import}}^{US\,exposure}_{bp,t-2}$	9.450	28.331	12.119	10.136	6.855
	(0.739)***	(4.101)***	(0.823)***	(0.839)***	(1.449)***
$\widehat{LowProp_p \times \text{Import}}^{US\,exposure}_{bp,t-2}$		-20.330			
		(4.176)***			
$\widehat{LowVisibility_p \times \text{Import}}^{US\,exposure}_{bp,t-2}$			-13.081		
			(1.801)***		
F-stat first stage	977.9	773.1	978.0	890.7	503.0
SW F – interaction		157.9	1020.6		
SW F		73.0	1154.2		
Stock-Yogo 10% critical value		7.03	7.03		
Stock-Yogo 15% critical value		4.58	4.58		
Observations	97,499,954	64,299,322	97,499,954	92,579,497	4,917,261
Clusters	2,187,612	1,373,213	2,187,612	2,113,249	115,253
Mean dependent variable	0.03	0.03	0.02	0.02	0.09
bp, bt, f FE	Yes	Yes	Yes	Yes	Yes

Notes: Robust standard errors, adjusted for clustering by retailer-product, are in parentheses. The independent variables are standardized. Regressions control for neighborhood×product, neighborhood×time, and retailer fixed-effects. Percentage mean import probabilities are reported. The value of the Sanderson and Windmeijer (2016) conditional first-stage F-statistics (SW F) for the validity of the instruments is reported in columns (2) and (3), along with the Stock-Yogo critical values for a perfectly identified model with two endogenous variables. Appendix F.5 presents details on the sample used in each regression. Data is quarterly and spans 2015-2019.

given the perceived local demand. Similarly, if people in an area gain insights about a product being inconsistent with local taste or not popular among their network, this information might be transmitted to retailers, making them less likely to import it compared to retailers in other locations who did not receive this insight. Results relying on low visibility goods as a proxy of low propagation goods are consistent with this narrative, and reported in column (3). Consistent with the idea that retailers now serve local markets, local projections show that individual imports of a product decrease once retailers start selling it domestically (Appendix F.6).

Appendix I.3 provides an additional test in support of firms learning about the level of local demand by leveraging the imperfect overlap between employer-employee networks and the residential location of employees. The idea behind this exercise is that employees can be exposed to foreign products in their neighborhoods and transmit information about the existence of these products to their employers, which would be relevant under a product-discovery story. However, if employees reside in areas far away from the retailer's catchment area, outside its gravity zone, their insights should be less informative about the particular level of the local demand that their employer would face. Indeed, we find that retail firms do not respond to the exposure faced by employees who live far away from their catchment area, highlighting the relevance of local demand knowledge.

We further investigate the mechanism behind the result in column (1), by exploring which retailers are more likely to respond. Comparing columns (4) and (5), we find that supply effects are mainly driven by small retailers.[51] This is consistent with the notion that small retailers, due to lower productivity, are less able to pay the search costs associated with identifying new foreign products, which larger retailers can afford. Thus, small retailers are more likely to take advantage of information from consumers when choosing which products to import. In fact, as shown in column (2) of Table H.3, the effect is much larger for premium, more expensive products, for which importing for the first time without a notion of local demand might entail a greater risk. In addition, small retailers have a more direct connection with local consumers and might be more responsive to their requests and needs when choosing which products to source.[52] While Table 6 considers *all* retailers, and therefore relies on a proxy of each gravity zone (see Appendix I.1), Table I.4 presents estimations with the subsample for which customers' locations are available. The results and the narrative are consistent between tables—in fact, most results are statistically equal.

Taken together, our empirical investigation then reveals that retail firms learn from their local consumers about which products to source from abroad. This channel is particularly relevant for small retailers, who take advantage of the revealed local demand to identify the "preferred" varieties that align with local customer taste.

5.2 Survey Evidence on the Mechanism

We conduct a nationally representative survey spanning 700 retail firms. Given Costa Rica's size and number of retailers, this is a large-scale survey reaching about 4% of re-

[51] We define a retailer as small if its annual employees is below 30, which is below the 95th percentile of firm size—the median of this sample is only two employees, so it would represent only micro enterprises.

[52] For instance, if consumers frequent a store inquiring about a product, a small retailer whose manager is at the shop might be more likely to react to these inquiries.

tail firms.[53] Respondents were required to have a relatively deep knowledge about the business.[54] The large-scale survey was conducted in partnership with CID-Gallup.[55] All interviewers underwent a one-week training on survey administration, and the instrument was refined through an initial pilot phase. The survey was conducted primarily by telephone during the second and third quarters of 2024; in exceptional cases, face-to-face interviews were carried out when no alternative data collection methods were feasible. Survey responses were linked to administrative firm characteristics, so that we can explore how responses vary across, for instance, firm sizes.

The survey was designed to examine the mechanisms previously documented. Retailers were asked about the role of their customers in determining their choice of which goods to import. We also designed the survey to investigate if results in columns (4) and (5) of Table 6 reflect actual experiences of retailers, and to understand *how* would a small retailer learn about what its customers import and if they liked the item; i.e., the mechanism behind responses in columns (2) and (3). The survey also aimed to confirm the timing of equation (10) was reasonable, and to acquire a survey-based counterpart of the results based on employees outside the retailer's gravity zone presented in Appendix I.3. We report the questions asked in Appendix I.4.

Results The survey results align the real-world experiences of retailers with our empirical findings. Starting with the broader result, we document that 81% of respondents reported receiving customer feedback on product selection, underscoring the influence of consumer preferences on inventory decisions. Additionally, aligned with column (1) of Table 6, 60% of retailers noted that observing potential customers buying new products from abroad would make them more likely to start importing and selling those products locally. This suggests a dynamic local market adaptation to global consumer behaviors informed by direct consumer imports.

Furthermore, consistent with columns (4) and (5) of Table 6 and as reported in panel (a) of Figure I.2, the survey documents that small retailers rely more heavily on direct customer imports to gauge the local demand for a potential new product than large firms. Aligned with the empirical analysis, these results indicate a distinct

[53] This leads to results with a 95% confidence interval.

[54] Respondents were categorized based on their roles within the company, as owner, manager or director, employee involved with sales, or employee in a capacity unrelated to sales.

[55] CID-Gallup has over than 40 years conducting market research and surveys in Latin America. It was established in 1977 in Costa Rica.

strategic approach based on company size, with smaller domestic firms being more responsive to local consumer needs via our mechanism.

We next leverage the survey to delve into understanding the mechanism behind the results. Columns (2) and (3) of Table 6 indicate that retailers are more responsive to products that propagate widely among individuals; but how do retailers learn about the products consumers are importing and whether they like them? Panel (b) of Figure I.2 illustrates the responses. Around half of the retailers gain insights into which imported products interest their customers when people physically visit the store and ask about the availability of these products. This primary channel for gathering information is followed by social media, local market studies, and consultations with family and friends. Also related to the mechanism, and in line with the results in Appendix I.3, retailers are over five times more likely to consult employees who reside near the store than employees who reside further away. In summary, this survey evidence supports our empirical analysis and reinforces the narrative that retail firms, especially smaller ones, rely on local consumers' importing experiences to inform their decisions on which products to source. The key channel to gain these insights is direct customer inquiries made in their physical stores.

5.3 Total Effect and Distributional Implications

We now bring together the effects we documented, both in terms of direct demand externalities across individuals and indirect demand effects from customers to retailers, to understand their total impact on local demand. To do so, let's recall that in Section 4.6 we estimated the impact of an increase in exposure from the U.S. on *individual imports*. Similarly, the back-of-the-envelope approximation of the increase in retailers' imports of a particular good is given by the product of:

$$C = \underbrace{\Delta Prob(Import_f)}_{\beta_3 \Delta ShareImporters^{US,direct}} \times Retailers \times AvgValue^{Retailer}, \tag{11}$$

where *Retailers* is the total number of retailers and $AvgValue^{Retailer}$ is the average value of a retailer's shipment (\$4,152). Then, equation (11) along with equation (7) imply that local retailers' imports of a product would increase in \$159 for every dollar exogenously imported by an individual directly connected with the U.S. (i.e., C/A).

What is the total multiplier given the demand propagation from the U.S., to individuals with relatives, to others in their network, and finally to retailers? To illustrate, suppose $\Delta\widetilde{E}_{bpt}$ corresponds to a \$100 increase in U.S. per capita spending on a product. Then, Costa Rican imports of this product per individual with a relative abroad would increase in approximately \$5, given the degree of interconnectedness between the two countries and the strength of the direct and indirect demand propagation.[56]

This is a total sizable effect, which can be decomposed into additional imports due to the direct externality and individuals' responses—8% of the effect—and the indirect externality and imports of retail domestic firms—the remaining 92% of the effect. These magnitudes underscore how accounting for this new supply-side indirect externality is key when estimating a full response.

Finally, the documented externalities have distributional implications. First, the diffusion channel may enable lower-income families to benefit indirectly from migration. Those with relatives abroad have, on average, 12% higher incomes, and consequently, are more likely to import in response to foreign demand shocks.[57] Second, those who import directly have, on average, 49% higher income than those who do not. As retailers learn from consumer imports and subsequently introduce products, the indirect externality leads to new product varieties available locally for lower-income households. These forces indicate that diffusion of products via retailers can lead to variety gains which are more evenly distributed across income groups.

6 Concluding Remarks

This paper investigates the role of direct and indirect externalities in propagating demand shocks both across and within countries—among individuals and from individuals to retailers. Our analysis makes three key contributions. First, we demonstrate how product-specific demand shocks propagate through international migrant networks, suggesting that migration policies can influence global product diffusion. Second, we quantify how an individual's first import increases the probability that others in her network will also import the same product—a force that varies across

[56]Like in Section 4.6, this multiplier is in per capita terms, where the denominator in Costa Rica is the number of individuals with a relative abroad, as these individuals act as the bridge through which demand propagates. A more connected set of countries (with more directly exposed individuals) would face an overall stronger propagation.

[57]Incomes are calculated based on wage income from social security records.

network types and product characteristics. Third, we use this propagation across individuals as a building block to show that retailers learn about local demand for foreign products by observing consumer imports; when consumer imports signal robust local demand, retailers respond by importing the product themselves.

A decomposition of these effects reveals that while individual responses (direct externalities) matter, retailer responses (indirect externalities) account for 92% of the overall impact, underscoring their central role in expanding access to product varieties. Notably, although direct imports are concentrated among higher-income households, the diffusion of product varieties via retail channels ultimately benefits lower-income families.

The presence of these externalities implies that gains from trade may be larger than previously documented, generating a multiplier effect for policies that stimulate foreign product demand, such as lower tariffs or relaxed import requirements. This insight is especially relevant for developing markets and is central to debates over tariff exemptions for individual imports. For example, in the U.S., the "de minimis" provision of Section 321 of the Tariff Act of 1930—which waives tariffs for shipments valued under $800—is currently under review by several coalitions.

Finally, this paper is the first to study the determinants of individual imports. Although historically uncommon, the rapid expansion of the direct-to-consumer market is only expected to accelerate due to increased internet penetration, improved logistics, and globalization. For instance, Temu, a Chinese app that allows for foreign direct-to-consumer purchases, was Apple's most downloaded free app in the U.S. for 2023, and low-value imports represented about 15% of all imports from China in 2021 according to the U.S. Customs Border Protection. Thus, this paper serves as an initial contribution to what promises to be fertile ground for future research.

References

Acosta, M., Cox, L., 2019. The Regressive Nature of the U.S. Tariff Code: Origins and Implications. Technical Report. Working Paper, Columbia University.

Acosta, P., Calderon, C., Fajnzylber, P., Lopez, H., 2008. What is the impact of international remittances on poverty and inequality in Latin America? World Development 36, 89–114.

Agarwal, A., Singh, R., Toshniwal, D., 2018. Geospatial sentiment analysis using twitter data for UK-EU referendum. Journal of Information and Optimization Sciences 39, 303–317.

Akbarpour, M., Malladi, S., Saberi, A., 2023. Just a Few Seeds More: The Inflated Value of Network Data for Diffusion. Working Paper.

Allen, T., 2014. Information Frictions in Trade. Econometrica 82, 2041–2083.

Alvarez, F., Argente, D., Lippi, F., Méndez, E., Van Patten, D., 2023. Strategic Complementarities in a Dynamic Model of Technology Adoption: P2P Digital Payments.

Aronow, P.M., Samii, C., 2017. Estimating average causal effects under general interference, with application to a social network experiment .

Atkin, D., Faber, B., Gonzalez-Navarro, M., 2018. Retail Globalization and Household Welfare: Evidence from Mexico. Journal of Political Economy 126, 1–73.

Bai, J., Chen, M., Liu, J., Mu, X., Xu, D.Y., 2020. Search and information frictions on global e-commerce platforms: Evidence from Aliexpress. Technical Report. National Bureau of Economic Research.

Bailey, M., Johnston, D.M., Kuchler, T., Stroebel, J., Wong, A., 2022. Peer Effects in Product Adoption. American Economic Journal: Applied Economics .

Bandiera, O., Barankay, I., Rasul, I., 2009. Social Connections and Incentives in the Workplace: Evidence From Personnel Data. Econometrica 77, 1047–1094.

Batch, A., Bridgman, B., Dunn, A., Gholizadeh, M., 2024. Consumption zones. Journal of Economic Geography , lbae035URL: https://doi.org/10.1093/jeg/lbae035, doi:10.1093/jeg/lbae035.

Beine, M., Docquier, F., Rapoport, H., 2008. Brain Drain and Human Capital Formation in Developing Countries: Winners and Losers. The Economic Journal 118, 631–652.

Borusyak, K., Hull, P., 2020. Non-Random Exposure to Exogenous Shocks: Theory and Applications. Working Paper 27845. National Bureau of Economic Research. URL: http://www.nber.org/papers/w27845, doi:10.3386/w27845.

Borusyak, K., Hull, P., 2023. Nonrandom Exposure to Exogenous Shocks. Econometrica 91, 2155–2185.

Borusyak, K., Jaravel, X., 2021. The Distributional Effects of Trade: Theory and Evidence from the United States. Technical Report. National Bureau of Economic Research.

Brock, W.A., Durlauf, S.N., 2001. Discrete Choice with Social Interactions. The Review of Economic Studies 68, 235–260.

Broda, C., Weinstein, D.E., 2006. Globalization and the Gains from Variety. The Quarterly Journal of Economics 121, 541–585.

Caplin, A., Leahy, J., 1998. Miracle on Sixth Avenue: Information Externalities and Search. The Economic Journal 108, 60–74.

Chaney, T., 2014. The Network Structure of International Trade. American Economic Review 104, 3600–3634.

Charles, K.K., Hurst, E., Roussanov, N., 2009. Conspicuous Consumption and Race. The Quarterly Journal of Economics 124, 425–467.

Conley, T.G., Udry, C.R., 2010. Learning about a New Technology: Pineapple in Ghana. American Economic Review 100, 35–69.

Conte, M., Cotterlaz, P., Mayer, T., 2022. The CEPII Gravity Database .

Couture, V., Faber, B., Gu, Y., Liu, L., 2021. Connecting the Countryside via E-Commerce: Evidence from China. American Economic Review: Insights 3, 35–50.

De Giorgi, G., Frederiksen, A., Pistaferri, L., 2019. Consumption Network Effects. The Review of Economic Studies 87, 130–163. doi:10.1093/restud/rdz026.

De Giorgi, G., Pellizzari, M., Redaelli, S., 2010. Identification of Social Interactions through Partially Overlapping Peer Groups. American Economic Journal: Applied Economics 2, 241–75.

EBANX, B.B..., 2023. Digital payments connecting businesses and people in rising economies: an overview of online commerce in Latin America and Africa, 2023. Online report. Accessed: 2023-09-13.

Faber, B., Fally, T., 2022. Firm Heterogeneity in Consumption Baskets: Evidence from Home and Store Scanner Data. The Review of Economic Studies 89, 1420–1459.

Fajgelbaum, P.D., Khandelwal, A., 2024. The Value of De Minimis Imports. Technical Report. National Bureau of Economic Research.

Fernandes, A.P., Tang, H., 2014. Learning to export from neighbors. Journal of International Economics 94, 67–84. URL: https://www.sciencedirect.com/science/article/pii/S0022199614000865, doi:https://doi.org/10.1016/j.jinteco.2014.06.003.

Furman, J., Russ, K., Shambaugh, J., 2017. US tariffs are an arbitrary and regressive tax. VoxEU blog .

Goldsmith-Pinkham, P., Imbens, G.W., 2013. Social Networks and the Identification of Peer Effects. Journal of Business & Economic Statistics 31, 253–264.

Head, K., Mayer, T., 2014. Gravity Equations: Workhorse,Toolkit, and Cookbook, in: Handbook of international economics. Elsevier. volume 4, pp. 131–195.

Hottman, C.J., Monarch, R., 2020. A matter of taste: Estimating import price inflation across U.S. income groups. Journal of International Economics 127, 103382.

Im, K.S., Pesaran, M.H., Shin, Y., 2003. Testing for unit roots in heterogeneous panels. Journal of econometrics 115, 53–74.

Jordà, Ò., Schularick, M., Taylor, A.M., 2020. The effects of quasi-random monetary experiments. Journal of Monetary Economics 112, 22–40.

Jordà, , 2005. Estimation and Inference of Impulse Responses by Local Projections. American Economic Review 95, 161–182.

Juhász, R., Steinwender, C., 2018. Spinning the web: Codifiability, information frictions and trade. NBER working paper 18.

Kuhn, K., Galloway, T., Collins-Williams, M., 2016. Near, far, and online: Small business owners' advice-seeking from peers. Journal of Small Business and Enterprise Development .

Levin, A., Lin, C.F., Chu, C.S.J., 2002. Unit root tests in panel data: asymptotic and finite-sample properties. Journal of econometrics 108, 1–24.

Manski, C.F., 1993. Identification of Endogenous Social Effects: The Reflection Problem. The Review of Economic Studies 60, 531–542.

Maurer, J., Meier, A., 2008. Smooth it Like the 'Joneses'? Estimating Peer-Group Effects in Intertemporal Consumption Choice. The Economic Journal 118, 454–476.

McCully, B.A., Jaccard, T., Albert, C., 2024. Immigrants, imports, and welfare: Evidence from household purchase data. Technical Report. RF Berlin-CReAM Discussion Paper Series.

Méndez, E., Van Patten, D., 2022. Voting on a Trade Agreement: Firm Networks and Attitudes Toward Openness. Technical Report. National Bureau of Economic Research.

Moffitt, R.A., 2000. Welfare Benefits and Female Headship in U.S. Time Series. American Economic Review 90, 373–377.

Montiel Olea, J.L., Plagborg-Møller, M., 2021. Local Projection Inference Is Simpler and More Robust Than You Think. Econometrica 89, 1789–1823.

Reardon, T., Timmer, C.P., Barrett, C.B., Berdegué, J., 2003. The Rise of Supermarkets in Africa, Asia, and Latin America. American Journal of Agricultural Economics 85, 1140–1146.

Sanderson, E., Windmeijer, F., 2016. A weak instrument F-test in linear IV models with multiple endogenous variables. Journal of econometrics 190, 212–221.

Startz, M., 2016. The value of face-to-face: Search and contracting problems in Nigerian trade. Available at SSRN 3096685 .

Steinwender, C., 2018. Real effects of information frictions: When the states and the kingdom became united. American Economic Review 108, 657–696.

Wooldridge, J.M., 2002. Econometric Analysis of Cross Section and Panel Data. MIT Press. Cambridge, MA.

Online Appendix for

Cross-Border Product Adoption: Individual Imports, Migrant Networks, and Domestic Retailers

March 19th, 2025

David Argente
Yale University and NBER

Esteban Méndez
Central Bank of Costa Rica

Diana Van Patten
Yale University and NBER

Table of Contents

A Details on New Stylized Facts

Table A.1: Top HS-4 Codes Imported by Individuals

Code	**Description**	**Share**
6204	Women's or girls' suits, ensembles, jackets, dresses, skirts, divided skirts, trousers, bib and brace overalls, breeches, and shorts.	10%
8703	Motor cars and other motor vehicles designed to transport people (other than those of heading 8702, including station wagons and racing cars).	5%
6206	Women's or girls' blouses, shirts, and shirt-blouses.	2%
4202	Trunks, suitcases, vanity cases, executive cases, briefcases, school satchels, and similar containers; traveling bags, backpacks, handbags, and similar products.	2%
8708	Parts and accessories of the motor vehicles of headings 8701 to 8705.	2%
6205	Men's or boys' shirts.	2%
6110	Sweaters, pullovers, sweatshirts, waistcoats (vests), and similar articles, knitted or crocheted.	2%
3926	Articles of plastics and articles of other materials of headings 3901 to 3914 (includes plastic parts or accessories, but can vary greatly).	2%
9503	Tricycles, scooters, pedal cars, and similar wheeled toys; dolls' carriages; dolls and other toys; reduced-size ("scale") models and similar recreational models.	2%
6302	Bed linen, table linen, toilet linen, and kitchen linen.	2%

Notes: The table documents the top HS-4 codes imported by individuals in Costa Rica. This ranking results from collapsing imports from HS-10 to HS-4 categories to be more informative; otherwise most top 10 imports would belong to the 6204 category. Data spans 2015-2019.

Table A.2: Top HS-4 Codes for Imported Final Goods: All and by Retailers

All Final Goods		Retailers' Imports	
(1)	(2)	(3)	(4)
Code	**Share**	**Code**	**Share**
8536	3%	8708	4%
7318	3%	6204*	3%
3926*	2%	3926*	2%
3923	2%	8536	2%
4016	2%	7318	2%
8544	2%	4202*	2%
7326	1%	8421	2%
8481	1%	4016	1%
8302	1%	8481	1%
8482	1%	6206*	1%

Notes: The table documents the top HS-4 codes for *all* final goods imported in Costa Rica (column (1)) and for final goods imported by retail firms (column (3)), with the respective shares. This ranking results from collapsing imports from HS-10 to HS-4 categories and does not weight by product value, and rather reports the most commonly imported goods. We denote with an asterisk the top codes which coincide with the top categories imported by individuals and reported in Table A.1. Data spans 2015-2019.

Figure A.1: Imports by HS-10 Code by Individuals

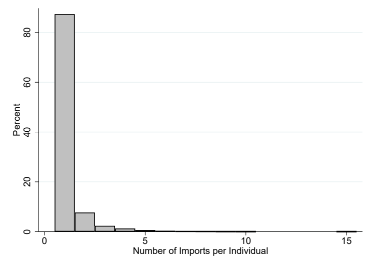

Notes: The histogram shows the total number of imports by HS-10 code by individual. As shown, most product categories are only imported by each person once.

B Conceptual Framework

In what follows, we describe a simple framework to think about demand externalities in the adoption of foreign products, both from an individual's perspective and from the point of view of the firm.

Setup N consumers in a network (e.g., neighborhood) want to buy a variety in a product category; they can either buy a domestic variety with a known payoff $D > 0$ or buy a foreign variety abroad with an uncertain (potentially greater) payoff. In a given period, each consumer decides whether to buy domestically and collect D or to search (online) for a product variety abroad. Each consumer who searches finds a variety abroad and decides whether to import it or not. Imported varieties cannot be returned; consumers would rather wait than purchase the wrong variety.[58]

We assume that consumers are risk neutral and maximize utility discounting future periods by $\rho \in (0, 1)$. The utility generated by an imported variety depends on $x\theta$, where x is consumer-specific and θ is common to all consumers. Consumers are ex-ante identical and have identical priors concerning the distributions of x and θ. We assume that x is drawn independently from a uniform distribution on $[0, 1]$ and is revealed to consumers when they find a variety (before they purchase it). All consumers know that θ is distributed uniformly on $[0, 2]$, but information on θ is revealed only after a consumer has imported the product, when both x and θ become public, and before the next period begins. This is, the first buyers do not observe θ, so their expected utility, if they decide to import a product, is $x\mathbb{E}(\theta) = x$. We assume that, after a variety is imported for the first time, θ becomes public and subsequent

[58]This assumption simplifies the analysis and it is reasonable in the case of Costa Rica, where the costs of returning an item are often too high; anecdotally, consumers often absorb the cost of internationally shipping back the item, the cost of processing the return, and suffer the delays of international shipping. In fact, in the data, *only 0.01% of individual imports are returned.*

buyers get $x\theta$. Thus, consumers want to learn about θ from others.

There is a single firm in the network, which sells to all consumers who buy the variety domestically. The firm also learns θ from the initial importers. If consumers decide to search for a variety abroad, the firm can choose to pay a fixed cost to make a once-and-for all decision to become an importer and sell the product variety domestically. Our information assumptions have two phases: uninformed and informed.

Uninformed Phase The uninformed phase takes place before any consumer has imported the foreign variety (i.e., before there is public information on θ).[59] In this phase, consumers search online and simply decide whether to import or not based on their draw of x. Intuitively, if x is sufficiently large, consumers decide to import the product variety. Formally, the uninformed agent maximizes

$$V_U(x) = \max\{x, \rho\,[p\,\mathbb{E}V_U + (1-p)\,\mathbb{E}V_I]\}, \tag{12}$$

where p is the endogenous probability that an agent remains uninformed (i.e., all other agents do not import) and $\mathbb{E}V_U$ and $\mathbb{E}V_I$ denote the expected value of being uninformed and informed, respectively. Thus, the consumer imports if $x \geq \widehat{x}$, where $\widehat{x} = \rho\,[p\mathbb{E}V_U + (1-p)\mathbb{E}V_I]$.

Informed Phase The informed phase begins after the first cohort has imported the product variety and θ becomes public. Consumers make two decisions. First, given θ, they decide whether to buy domestically (collect D) or continue searching online. Second, if they search, they must decide whether to import or not. Thus, their strategies determine a set of values $\theta \geq \bar{\theta}$ that warrant continued search, and a set of qualities $x \geq \bar{x}$ that determine whether the consumers import the variety or not. The value of an optimal strategy for an informed agent who decides to search is

$$V_I(x,\theta) = \max\{x\theta, \rho\mathbb{E}[V_I(x',\theta)]\}, \tag{13}$$

where the first term is the value of importing the product variety and the second term is the discounted value of keep searching. The cutoff $\bar{\theta}$ is pinned down where the value of searching equals the payoff of buying the variety domestically (i.e., $\mathbb{E}[V_I(x,\bar{\theta})] = D$). In turn, the cutoff $\bar{x}(\theta)$ is the value of x that makes consumers indifferent between importing or keep searching (i.e., $\bar{x}(\theta) = \frac{1}{\theta}\rho\mathbb{E}[V_I(x',\theta)]$).

Also in this phase, the firm learns θ. If consumers buy domestically (i.e., $\theta < \bar{\theta}$), the firm sells them the domestic variety. If consumers search (i.e., $\theta \geq \bar{\theta}$), the firm

[59] We assume that D is small enough so that consumers decide to search for imported products when θ is unknown. Thus, there are no domestic sales in the uninformed phase. In Section B.1, we show the upper bound of D that is a sufficient condition for search.

can decide to pay a fixed cost c and import the good. The firm's optimal strategy is

$$W_I(\theta) = \max\{0, -c + \rho\mathbb{E}[V_I(x', \theta)]\widetilde{N}\}, \tag{14}$$

where, without loss of generality, we normalize firm's profits to zero when it chooses not to import. If the firm decides to import the product, consumers' outside option is to buy abroad. Thus, under Bertrand competition, the firm sets the highest possible price that still prevents consumers from buying abroad. As a result, the firm sells to all other consumers who have not yet bought the good abroad $\widetilde{N} \equiv N\bar{x}(\theta)\widehat{x}$ and obtains all surplus. The firm imports if $\theta \geq \widetilde{\theta}$; the cutoff is pinned down at the point where the cost and the expected gains of importing are equal, $c = \rho\mathbb{E}[V_I(x, \widetilde{\theta})]N\bar{x}(\theta)\widehat{x}$.

Equilibrium We focus on equilibria in which the decision rules depend only on information that is payoff-relevant. Further, since all consumers are ex-ante identical, we look for a symmetric Nash Equilibrium. The stationary equilibrium involves a set of cutoff rules, summarized below.

Definition 1. An equilibrium consists of cutoffs $\bar{\theta} \in [0, 2]$, $\widetilde{\theta} \in [0, 2]$, $\widehat{x} \in [0, 1]$, and a function $\bar{x}(\bar{\theta}) : [\bar{\theta}, 1] \to [0, 1]$ such that the following strategy is optimal: (a) in the uninformed phase, only varieties with $x \geq \widehat{x}$ are imported, (b) in the informed phase, consumers search happens only if $\theta \geq \bar{\theta}$ and consumers import varieties with $x \geq \bar{x}$, and (c) in the informed phase, the firm imports the good to sell domestically if $\theta \geq \widetilde{\theta}$.

Properties of the Equilibrium Proposition 4 in Appendix B.1.3 establishes both existence and uniqueness of an equilibrium of the form given by Definition 1. The equilibrium has intuitive properties, which we describe in a set of propositions, each followed by its intuition. The proofs of all propositions can be found in Appendix B.2.

Proposition 1. The equilibrium level of $\widehat{x}$ is increasing in D.

In the uninformed phase, consumers are less willing to search online (or import) a variety if the value of the domestic option, D, is high. This result implies that a consumer would never choose to search for a product that is already available domestically with high enough quality or appeal.

Proposition 2. If $p > 0$, then $\widehat{x} > \bar{x}$ and $\mathbb{E}V_U < \mathbb{E}V_I$.

Consumers demand a higher x to import in the uninformed stage (i.e., $\widehat{x} > \bar{x}$), and the possibility of waiting for another consumer to import the good gives rise to a free rider problem. This result derives from an demand externality *that the model generates endogenously*: consumers do not internalize that importing a variety provides valuable information to other consumers in their network. Thus, the equilibrium is inefficient and there is a delay in the adoption of imported varieties, since the expected payoff in the informed phase is greater than that in the uninformed phase.

PROPOSITION 3. $\bar{\theta} \in (0,1)$ is increasing in D and $\widetilde{\theta} \in [\bar{\theta}, 2)$ is decreasing in N.

Lastly, in the informed phase, only varieties with high enough θ relative to the domestic option are imported, either by consumers or the firm. Note $\bar{\theta}$ is strictly less than $\mathbb{E}(\theta) = 1$, reflecting the value of information for consumers and for the firm. Further, the firm is more likely to import a variety and sell it domestically if its market size is large. Figure B.2 summarizes the model's solution for different values of x and θ.

From Model to Data Our empirical analysis is guided by this framework. In the uninformed phase, individuals decide whether to import independently from others; there are no common shocks or shared characteristics among peers. As the latter is unlikely to hold empirically, Section 4.1 proposes a strategy to leverage plausibly exogenous demand shocks to the likelihood of importing. Through the lens of the model, these shocks can be understood as shifters of x and can lead to importing varieties without good domestic alternatives (Proposition 1). In the informed phase, the model features direct externalities: once a person imports a variety, peers in her network may become more likely to import it as well (Proposition 2). Whether these externalities exist is an empirical question, which Section 4.5 explores for different networks. Finally, there is also an indirect externality: firms respond to the revealed information by importing, but only when the expected gains are sufficiently large (Proposition 3). In fact, firms only import varieties with strong enough propagation among consumers after they are imported. Section 5 analyzes retailers' responses after individuals import, differentiating between goods with strong and weak propagation.

B.1 Solution

The solution of the model follows closely Caplin and Leahy (1998). We solve the model backwards beginning with the informed phase. Given that the actual value of θ is known in this phase, consumers can compute the reservation value $\bar{x}$ by comparing the value of searching online and the value of importing the product variety. Similarly, consumers can compute $\bar{\theta}$ by comparing the value of buying domestically and the value of searching. Finally, knowing $\bar{x}$ and $\bar{\theta}$, the decision of whether to import the product variety or keep searching online in the uninformed phase pins down $\widehat{x}$. For the firm, the solution is simpler. Since there are no domestic sales in the uninformed phase, in the informed phase the firm decides whether to import the variety or not given demand and θ.

B.1.1 Informed Phase

Recall that in this phase consumers' maximize

$$V_I(x,\theta) = \max\{x\theta, \rho\mathbb{E}[V_I(x',\theta)]\}.$$

where it follows that the reservation level of x is pinned down by:

$$\begin{aligned}\bar{x}(\theta) &= \frac{1}{\theta}\rho\mathbb{E}[V_I(x',\theta)] \\ &= \rho\left(\frac{1+\bar{x}(\theta)^2}{2}\right)\end{aligned} \tag{15}$$

which shows that $\bar{x}(\theta) \equiv \bar{x}$.[60] Solving equation (15) focusing on solutions in the domain of $\bar{x}$, we find $\bar{x} = \frac{1-\sqrt{1-\rho^2}}{\rho}$.

Similarly, $\bar{\theta}$ is pinned down by:

$$\begin{aligned}D &= \mathbb{E}[V_I(x,\bar{\theta})] \\ &= \bar{\theta}\left(\frac{1+\bar{x}^2}{2}\right).\end{aligned} \tag{16}$$

Note that $\bar{\theta} < 1$ since this requires $D < \frac{1+\bar{x}^2}{2}$ and we know the upper bound $D < \frac{1+\widehat{x}^2}{2}$ from our assumption that initial search is more valuable than a domestic purchase in the uninformed phase. Letting $d \equiv \frac{D}{1+\bar{x}^2}$, we can write $\bar{\theta} = 2d$ where $d \in (0, \frac{1}{2})$.

Lastly, $\widetilde{\theta}$ is pinned down by:

$$\begin{aligned}c &= \rho\mathbb{E}[V_I(x,\theta)]\widetilde{N}\} \\ &= \rho\widetilde{\theta}\left(\frac{1+\bar{x}^2}{2}\right)\widetilde{N},\end{aligned} \tag{17}$$

where $\widetilde{N} \equiv N\bar{x}(\theta)\widehat{x}$. Using the definition of d, we can write $\widetilde{\theta} = \frac{2cd}{\rho\widetilde{N}D}$.

B.1.2 Uninformed Phase

In this phase, an uninformed consumer maximizes

$$V_U(x) = \max\{x, \rho\left[p\mathbb{E}V_U + (1-p)\mathbb{E}V_I\right]\},$$

[60] The second equality in equation (15) follows from $\mathbb{E}[\max\{x, x'|x' = \bar{x}\}] = \frac{1+\bar{x}^2}{2}$ when x and x' are two independent random variables uniformly distributed on $[0,1]$.

where $\widehat{x}$ satisfies the indifference condition between importing or searching and $p = \widehat{x}^{N-1}$ is the probability that an agent remains uninformed. Note that $\mathbb{E}V_U = \frac{1+\widehat{x}^2}{2}$ so that assuming $D < \frac{1+\widehat{x}^2}{2}$ is sufficient so that consumers decide to search for imported varieties when θ is unknown. The expected value of being informed in this phase is:

$$\begin{aligned}\mathbb{E}V_I &= \int \left[\max\left\{D, \int V_I(x,\theta)dx\right\}\right] d\theta \\ &= (1+d^2)\left(\frac{1+\bar{x}^2}{2}\right) \\ &= (1+d^2)\frac{\bar{x}}{\rho},\end{aligned} \tag{18}$$

where we use equation (15) in the last line to simplify equation (18). The reservation acceptance level $\widehat{x}$ can be found using the indifference condition

$$\begin{aligned}\widehat{x} &= \rho\left[p\mathbb{E}V_U + (1-p)\mathbb{E}V_I\right] \\ &= \rho\left[\widehat{x}^{N-1}\frac{1+\widehat{x}^2}{2} + (1-\widehat{x}^{N-1})(1+d^2)\left(\frac{1+\bar{x}^2}{2}\right)\right]\end{aligned}$$

and, using equation (15) to eliminate ρ, we find

$$\widehat{x} = \widehat{x}^{N-1}\bar{x}\left(\frac{1+\widehat{x}^2}{1+\bar{x}^2}\right) + (1-\widehat{x}^{N-1})\bar{x}(1+d^2). \tag{19}$$

The value of $\widehat{x}(N)$ as N increase is relevant since it determines the severity of the free rider problem. In particular, the limit of $\widehat{x}(N)$ as N increases is

$$\lim_{N\to\infty}\widehat{x}(N) = \begin{cases}\rho EV_I & \text{if } \rho EV_I \leq 1 \\ 1 & \text{if } \rho EV_I > 1\end{cases} \tag{20}$$

where the value of ρEV_I is given in equation (18). Intuitively, if the expected value of being informed is very large (i.e., $\rho EV_I > 1$), the free rider problem becomes very serious and the wait for the first import of a product variety can become arbitrarily long (i.e., $\widehat{x}(N) = 1$).[61]

B.1.3 Equilibrium

Proposition 4. Let $\widetilde{N} \equiv N\bar{x}\widehat{x}$ and $d \equiv \frac{D}{1+\bar{x}^2}$. Then for $N > 1$, $\widetilde{N} > 1$, $\rho \in (0,1)$, $D \in (0, \frac{D}{2d})$, $d \in (0, \frac{1}{2})$ and $c \in [\frac{\rho\bar{\theta}\widetilde{N}D}{2d}, \frac{\rho\widetilde{N}D}{d}]$, there exists a unique equilibrium of the form given in Definition 1 with $\bar{x} = \frac{1-\sqrt{1-\rho^2}}{\rho}$, $\bar{\theta} = 2d$, $\widetilde{\theta} = \frac{2cd}{\rho\widetilde{N}D}$, and $\widehat{x} \in (\bar{x}, \min[1, \bar{x}(1+d^2)])$ which is uniquely determined.

[61] If $\rho EV_I \leq 1$ then $\lim_{N\to\infty}\widehat{x}(N)^N = 0$. If $\rho EV_I > 1$ then $\lim_{N\to\infty}\widehat{x}(N)^N = \dfrac{[\rho EV_I - 1]}{\rho[EV_I - 1]} \in (0,1)$.

B.1.4 State Space

Figure B.2 summarizes the model's solution for different values of x and θ. The orange rectangle is the area of the state space for which it is optimal to import a variety in the uninformed phase. The size of the blue rectangle relative to the orange indicates that, in the informed phase, there is a wider range of values for which importing is optimal. This difference explains the initial delay in the adoption of a foreign variety and its subsequent adoption by the network after someone imports it. Finally, the green rectangle shows the values for which it is optimal for the firm to import the variety and sell it domestically. The importing decision for the firm depends on the common quality or appeal of the foreign variety and on the amount of consumers in the network willing to buy it, but who have not already imported it on their own.

Figure B.2: Model Solution and Properties of the Equilibrium

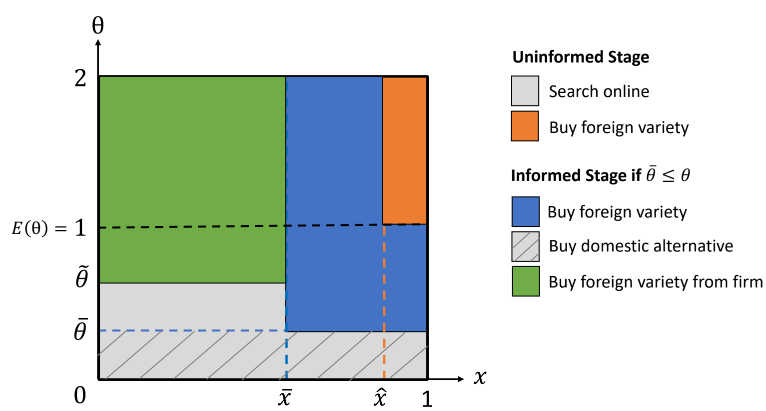

Notes: The figure shows the state space for θ and x, along with the equilibrium thresholds.

B.2 Proofs

Proof. (of Proposition 1) Let $D' > D$. Using equation (21), it can be shown that $H(\widehat{x}; D') < H(\widehat{x}; D)$. From Proposition 4 we know that $H(\widehat{x})$ is increasing in $\widehat{x}$ so an increase in D requires an increase in $\widehat{x}$ to restore equilibrium. Similarly, if $N' > N$ then $H(\widehat{x}; N') < H(\widehat{x}; N)$ if $\widehat{x} \leq (1+d^2)\bar{x}$. This condition can verified using equation (21) and evaluating it at $\widehat{x} = (1+d^2)\bar{x}$; in this case, $H(\widehat{x}) > 0$ which implies $\widehat{x} \leq (1+d^2)\bar{x}$. $\square$

Proof. (of Proposition 2) Our assumption that initial search is more valuable than a domestic purchase in the uninformed phase implies that $\frac{1+\widehat{x}^2}{2} > D = d(1+\bar{x}^2)$.

Thus, $\widehat{x} > \bar{x}$ follows from $d \in (0, \frac{1}{2})$. Moreover, using equation (21), we can verify that $H(\widehat{x}) > 0$ at $\widehat{x} = (1+d^2)\bar{x}$, so that $\widehat{x} \leq (1+d^2)\bar{x} = \rho \mathbb{E}V_I$. For $p > 0$ and $\widehat{x} > \bar{x}$, $\widehat{x} < \rho \mathbb{E}V_I$ since $\widehat{x} = \rho\left[p\mathbb{E}V_U + (1-p)\mathbb{E}V_I\right]$. Thus, $\mathbb{E}V_U < \mathbb{E}V_I$ $\square$

Proof. (of Proposition 3) From equation (16) we know that $\bar{\theta} = \frac{2D}{1+\bar{x}^2}$. Thus, $\bar{\theta} < 1$ since this requires $D < \frac{1+\bar{x}^2}{2}$ and we know $D < \frac{1+\widehat{x}^2}{2}$ from our assumption that initial search is more valuable than a domestic purchase in the uninformed phase. This implies that $d \in (0, \frac{1}{2})$ since $\bar{\theta} = 2d$ and using equation (16) it is straightforward to show that $\frac{\partial \bar{\theta}}{\partial D} > 0$. Moreover, note that $\widetilde{\theta}$ is bounded from below by $\bar{\theta}$, since for values of θ below $\bar{\theta}$ consumers prefer to purchase products available domestically. Using equation (17) it is easy to verify that $\frac{\partial \widetilde{\theta}}{\partial N} < 0$. $\square$

Proof. (of Proposition 4) To establish existence and uniqueness, we need to show that equation (19) provides a unique solution for $\widehat{x}$. Our assumption that initial search is more valuable than a domestic purchase in the uninformed phase implies that $\frac{1+\widehat{x}^2}{2} > D = d(1+\bar{x}^2)$. Since $d \in (0, \frac{1}{2})$, then $\widehat{x} > \bar{x}$. Thus, we need to show that equation (19) has a unique solution $\widehat{x} \in (\bar{x}, 1)$. We begin rewriting equation (19) as

$$H(\widehat{x}) = (1 + X_1^2)\widehat{x} + (X_0 - X_1)\widehat{x}^{N-1} - X_0 - X_1\widehat{x}^{N+1} \tag{21}$$

where $X_1 = \bar{x}$, $X_0 = (1+d^2)\bar{x}(1+\bar{x}^2)$ and $X_0 > X_1$. Note that $H(\widehat{x}) < 0$ if $\widehat{x} \to \bar{x}$ and $H(\widehat{x}) > 0$ if $\widehat{x} \to 1$. Thus, there exists a solution $\widehat{x} \in (\bar{x}, 1)$. Since $H(\widehat{x})$ starts below zero and ends above zero, we can show uniqueness by ruling out multiple zeros. This can be done by showing that the function is locally concave at any critical point. To do so, we first find the critical points

$$H'(\widehat{x}^*) = (1 + X_1^2) + (N-1)(X_0 - X_1)(\widehat{x}^*)^{N-2} - (N+1)X_1(\widehat{x}^*)^N = 0$$

and then we show that at any $\widehat{x}^* \in [0, 1]$, $H''(\widehat{x}^*) < 0$

$$\begin{aligned} H''(\widehat{x}^*) &= (N-2)(N-1)(X_0 - X_1)(\widehat{x}^*)^{N-3} - N(N+1)X_1(\widehat{x}^*)^{N-1} \\ &= \frac{1}{\widehat{x}^*}(N-2)(N-1)(X_0 - X_1)(\widehat{x}^*)^{N-2} - N(N+1)X_1(\widehat{x}^*)^N \\ &< \frac{N}{\widehat{x}^*}H'(\widehat{x}^*) = 0 \end{aligned}$$

Thus, $H(\widehat{x})$ only has one critical point. $\square$

C Setup: Additional Results

C.1 CEX vs. Other Expenditures Data

Relationship Between CEX and Imports Data In the U.S., many tradable products are imported. Thus, expenditure shares for these products in the CEX by region should co-move with the imports of these products in these areas. Following this idea, we use data on imports by customs districts in the U.S., adjusted using Freight Analysis Framework (FAF) data from the Department of Transportation as explained in Section 2, to assess the representativeness of the CEX at narrowly-defined categories and geographic areas. This notion follows Acosta and Cox (2019), who show that these customs districts data closely matches aggregate patterns in the CEX. Figure C.1 shows a strong correlation between expenditures in the CEX and expenditures based on customs districts data, when defining products as 4- or 6-digit HS codes, regions as Primary Sampling Units (PSUs), and time as quarters between 2015 and 2019. The correlation is strong both in levels as shown by Panels (A)-(B) and in changes as shown in Panels (C)-(D).

Figure C.1: Expenditure Shares in the CEX vs. Customs Districts

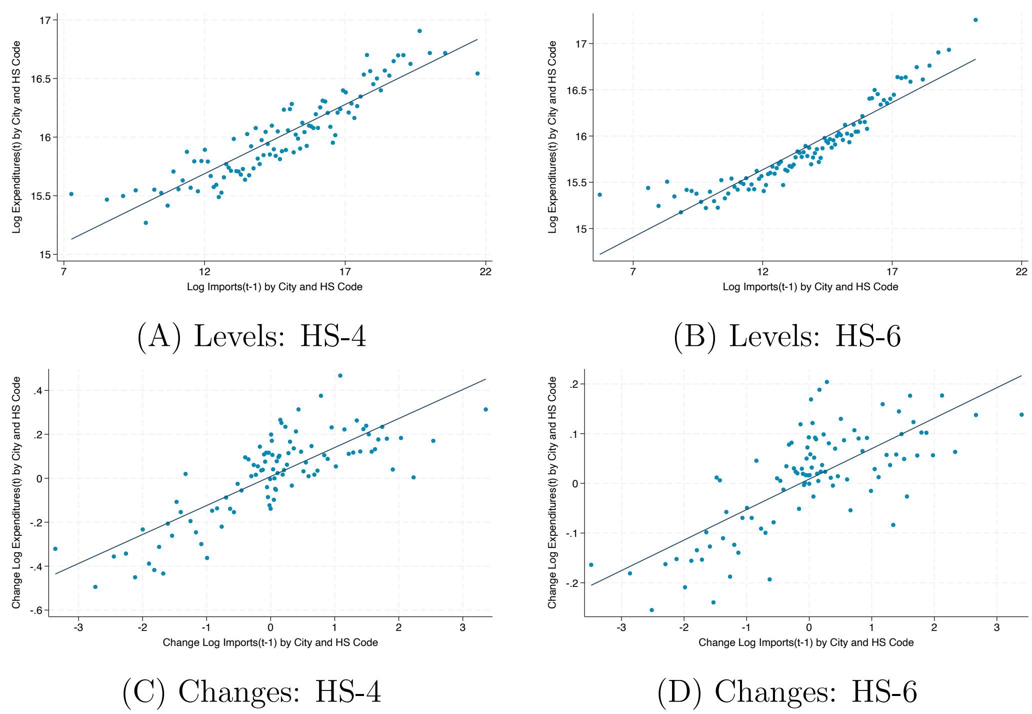

(A) Levels: HS-4

(B) Levels: HS-6

(C) Changes: HS-4

(D) Changes: HS-6

Notes: The figures shows the relation between expenditures in the CEX (vertical axis) and expenditure based on customs districts data (horizontal axis), when defining products as HS-4 or HS-6 product codes, regions as PSUs, and time as quarters for the period 2015-2019. Panels (A) and (B) show the correlation in levels for products as HS-4 and HS-6 codes, respectively. Panels (C) and (D) show the correlation for the same definition of products in changes, we trim the top and bottom one percent.

Relationship Between CEX and Debit Card Data We use data on debit card transactions by region and by type to further validate the CEX. This data comes from Facteus, a provider of financial data for business analytics. The data contain information on total expenditures by category at the zip-code level and with daily frequency. Approximately 10 million debit cards are included. The debit cards in the Facteus panel are issued by "challenger banks." The dataset spans from 2017 to 2019 and includes information of more than 200 Merchant Category Codes (MCCs), one per transaction, corresponding to the MCC standard by Visa and Mastercard. We manually create a bridge between MCC and Standard Industrial Classification (SIC) codes.[62] Figure C.2 shows a strong correlation between expenditures in the CEX and expenditures based on card transactions data, when defining products as HS-4 or HS-6 codes, regions as PSUs, and time as years. As in the case of customs data, the correlation is strong both in levels and in changes.

Figure C.2: Expenditure Shares in the CEX vs. Card Transactions

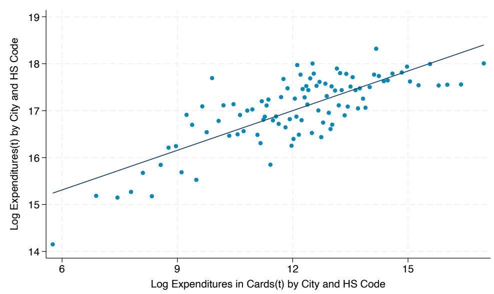

(A) Levels: HS-4

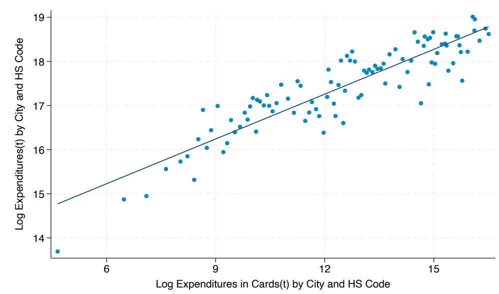

(B) Levels: HS-6

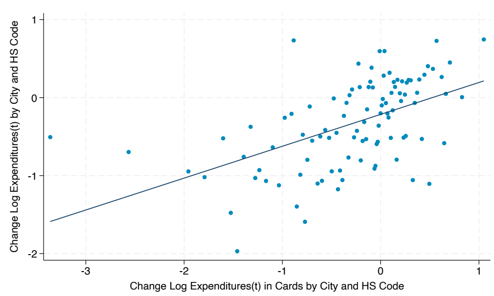

(C) Changes: HS-4

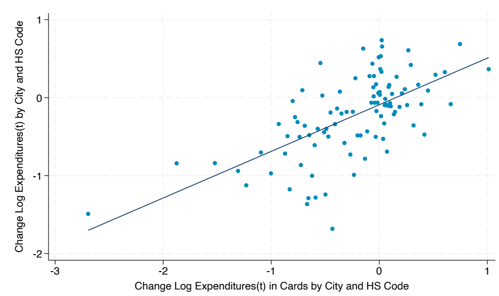

(D) Changes: HS-6

Notes: The figure shows the relation, both in levels and in changes, between expenditures based on the CEX (vertical axis) and on card transactions (horizontal axis), when defining products as HS-4 or HS-6 product codes, regions as PSUs, and time as years; we trim the top and bottom one percent.

[62]This bridge was created in parallel by two independent teams of RAs, then cross-checked, and finally revised by the authors. MCCs were derived from SIC codes, however, MCCs and SIC codes do not always correspond; in some cases, several SIC codes are consolidated into one MCC, while in other cases, such as for "T&E and direct marketing merchants," MCCs do not have corresponding SIC code.

C.2 Network Descriptive Statistics

Table C.1: Network Summary Statistics

Network type	Total number of networks (1)	Median individuals per network (2)
Neighbors	1,681	781
Coworkers	11,803	12
Friends	109,438	7

Notes: The table shows the total number of distinct networks per network type, along with the median number of people who compose each network. *These are only networks which have at least one person with a relative abroad.* See Appendix E for details on why network of friends are so numerous.

D Analysis of Exposure Measures

D.1 Variation and Serial Correlation Tests

Table D.2: Variation for Exposure Measures
(% of products with underlying variation at each HS-code level)

Panel (a): CEX				*Panel (b): U.S. imports*			
HS-4	HS-6	HS-8	HS-10	HS-4	HS-6	HS-8	HS-10
91.01	6.35	2.12	0.53	0	**59.95**	21.63	18.42

Notes: The table shows the percentage of products in our sample whose underlying variation is at each HS-code level, which tells us the level at which the exposure measures for our IV strategy are varying, depending on the source from which we obtain expenditures on each product code by region and time in the U.S. Panel (a) shows that most of the variation is at the HS-4 level when using the CEX. Panel (b) shows most of the variation is at the HS-6 level when relying on U.S. imports data by customs districts.

Serial Correlation Tests To assess the presence of serial correlation in $\widetilde{E}_{cpt}$, we conducted the Wooldridge test (Wooldridge, 2002) for autocorrelation in panel data. We find no evidence of serial correlation (i.e., the null hypothesis of no first-order autocorrelation could not be rejected). We also conducted unit root tests to determine the stationarity of $\widetilde{E}_{cpt}$. Specifically, we employed both the Levin-Lin-Chu test (Levin et al., 2002) and the Im-Pesaran-Shin (IPS) test (Im et al., 2003), two tests that are appropriate for panel data with a large number of cross-sections and a smaller number of time periods. The LLC test assumes a common unit root process, while the IPS test allows for individual unit root processes across panels. After allowing for several lags to account for possible autocorrelation and dynamic effects within the data, we corroborate the stationarity of $\widetilde{E}_{cpt}$. Table D.3 presents the results.

Table D.3: Serial Correlation Tests for $\widetilde{E}_{cpt}$

	Wooldridge Test (1)	Levin-Lin-Chu Test (2)	Im-Pesaran-Shin Test (3)
H_0	No serial autocorrelation	Panels contain unit roots	All panels contain unit roots
Criteria	Prob$>F = 0.2730$	One lag: P-val = 0.000 Two lags: P-val = 0.000	No lags: P-val = 0.000 Two lags: P-val = 0.000

Notes: The table shows different tests for serial correlation with their null hypothesis (H_0) and p-value, which corroborate the stationarity of $\widetilde{E}_{cpt}$. The first column shows results for the Wooldridge test (Wooldridge, 2002) for autocorrelation in panel data. We also employed the Levin-Lin-Chu test (Levin et al., 2002) and the Im-Pesaran-Shin (IPS) test (Im et al., 2003) in columns (2) and (3), respectively, two tests that are appropriate for panel data with a large number of cross-sections and a smaller number of time periods.

D.2 Variation in the Residuals

This section explores what are the drivers of the variation in the residuals estimated via equation (1), and in particular, whether they are driven by local entry and exit of brands within narrowly defined product categories.

Data To explore this question, we leverage microdata on the dynamics of different narrowly defined products, brands, and retailers across U.S. regions and time, for several product categories. Our source is *Consumer Insights* (CI), which provides consumer survey data for durable goods; items like home appliances, power tools, and electronics are all examples of goods that fit into this category.[63] This Durable IQ survey is run quarterly and surveys 600,000 households in the United States each year. The survey asks consumers a series of questions about the comprehensive purchases that they have made within the past 90 days, and is highly specific about features of each product. The data spans 2015-2022 and includes product type, brand, outlet (i.e., retailer, including Lowe's, Home Depot, or Walmart), online vs. in-store channel, demographics of users, and city where the purchase took place.

Strategy All the cities covered in the CEX are also available in the CI data. We then match product-brand pairs in the CI data to their corresponding categories in the CEX. In some cases products in both data sets coincide one-to-one. However, in other cases, CI may classify products differently from how they are described in which are part of the CEX category, or not include products related to some CEX categories.[64] Therefore, we focus on items that are available in both data sets and

[63]While the Nielsen IQ data would have provided information on more categories, its access is restricted whenever a co-author is not tenured or tenure-track faculty, a PhD student or postdoc. Thus, it is incompatible with using Costa Rican administrative data.

[64]For instance, a CEX category includes: car stereo, CB radio, mixer, speakers, stereo, amplifier, clock radio, receiver, turntable, walkie talkie, equalizer, compact disc players, short-wave radio, stereo system, tuner, satellite radio, audio cassette players/recorders. In

that roughly correspond with each other. Since our purpose is to understand what is driving the variation in the residuals using the detailed micro data available in CI, we also focus on CEX categories whose residualized expenditures significantly correlate with the residualized expenditures of the constructed product categories in the CI data.[65] Thus, we consider five main categories which include electronics, small appliances, major appliances, and tools.[66]

Results Column (1) of Table D.4 shows that the log expenditures in the CI data correlate well with those in the CEX after including all the battery of fixed effects; i.e., residualized expenditures in both datasets are correlated. Next, we explore the determinants of the variation in residualized expenditures. To do so, we define products in the CI data as product type-brand pairs; for instance, a JBL mini speaker. Column (2) shows that there is significant correlation between the residuals in expenditures in the CEX and changes in the number of product-brand pairs in the CI data; in fact, as shown on column (4), changes in available brands within product categories is the main driver of changes in expenditures in the CI data, with a correlation of 0.87. Columns (3) and (5) show the number of retailers in a location also correlates with changes in expenditures in both datasets, although not as salient as distinct product-brand pairs.

Table D.4: Correlation: Categories in the CEX vs. Consumer Insights (CI)

	Log Exp. (CEX)			Log Exp. (CI)	
	(1)	(2)	(3)	(4)	(5)
Log Exp. (CI)	0.3651***				
	(0.094)				
Log Products (CI)		0.3765**		0.8696***	
		(0.152)		(0.115)	
Log Retailers (CI)			0.2564*		0.5074***
			(0.147)		(0.097)
Observations	2,299	2,300	2,300	2,299	2,999
R-squared	0.673	0.671	0.671	0.983	0.979
sp, *st*, *pt* FE	Yes	Yes	Yes	Yes	Yes

Notes: The table displays the correlations between the *residuals* of (log) products and retailers and the (log) expenditures by product category in a given city in the CEX and CI data. Products are defined as product type-brand pairs. Robust standard errors are in parentheses. All regressions control for city×product, city×time, and product×time fixed-effects.

this case, the CI data has the following products which match the category in the CEX reasonably well: home theater in box, mini speaker, mini/shelf stereo system, separate receiver, soundbar, traditional speaker.

[65] Recall that we residualized CEX expenditures using product-city, city-time, and product-time fixed effects.

[66] Namely, the categories are: "Small electric kitchen appliances," "Stereos, radios, speakers, and sound components including those in vehicles," "Clothes washer or dryer (owned home)," "Portable heating and cooling equipment," and "Power tools."

We then zoom into the five main categories described above, and ask whether the dynamics of entry and exit of brands *within* narrowly defined products that below to each category can explain movements in the residualized expenditures. Figure D.3 plots these residuals across time across these categories and many different U.S. cities. First, in line with the serial autocorrelation test results in Table D.3, note that these residuals do not display from serial correlation. Second, in line with the results in Table D.4, the residualized expenditures (solid black lines) strongly co-move with the product type-brand residuals (dashed red line). In other words, the number of brands *within* narrowly defined products that below to each category seem to be driving changes in the residualized expenditures across regions. Finally, also in line with Table D.4, the dynamics of retailers also display important co-movement with the other residuals, albeit weaker than the one for product brands. Overall the analysis suggests that the changes in residuals are greatly driven by the differential entry and exit of products of different brands across space and time.

Figure D.3: Residuals per Product Category and Micro Data on Brand Dynamics

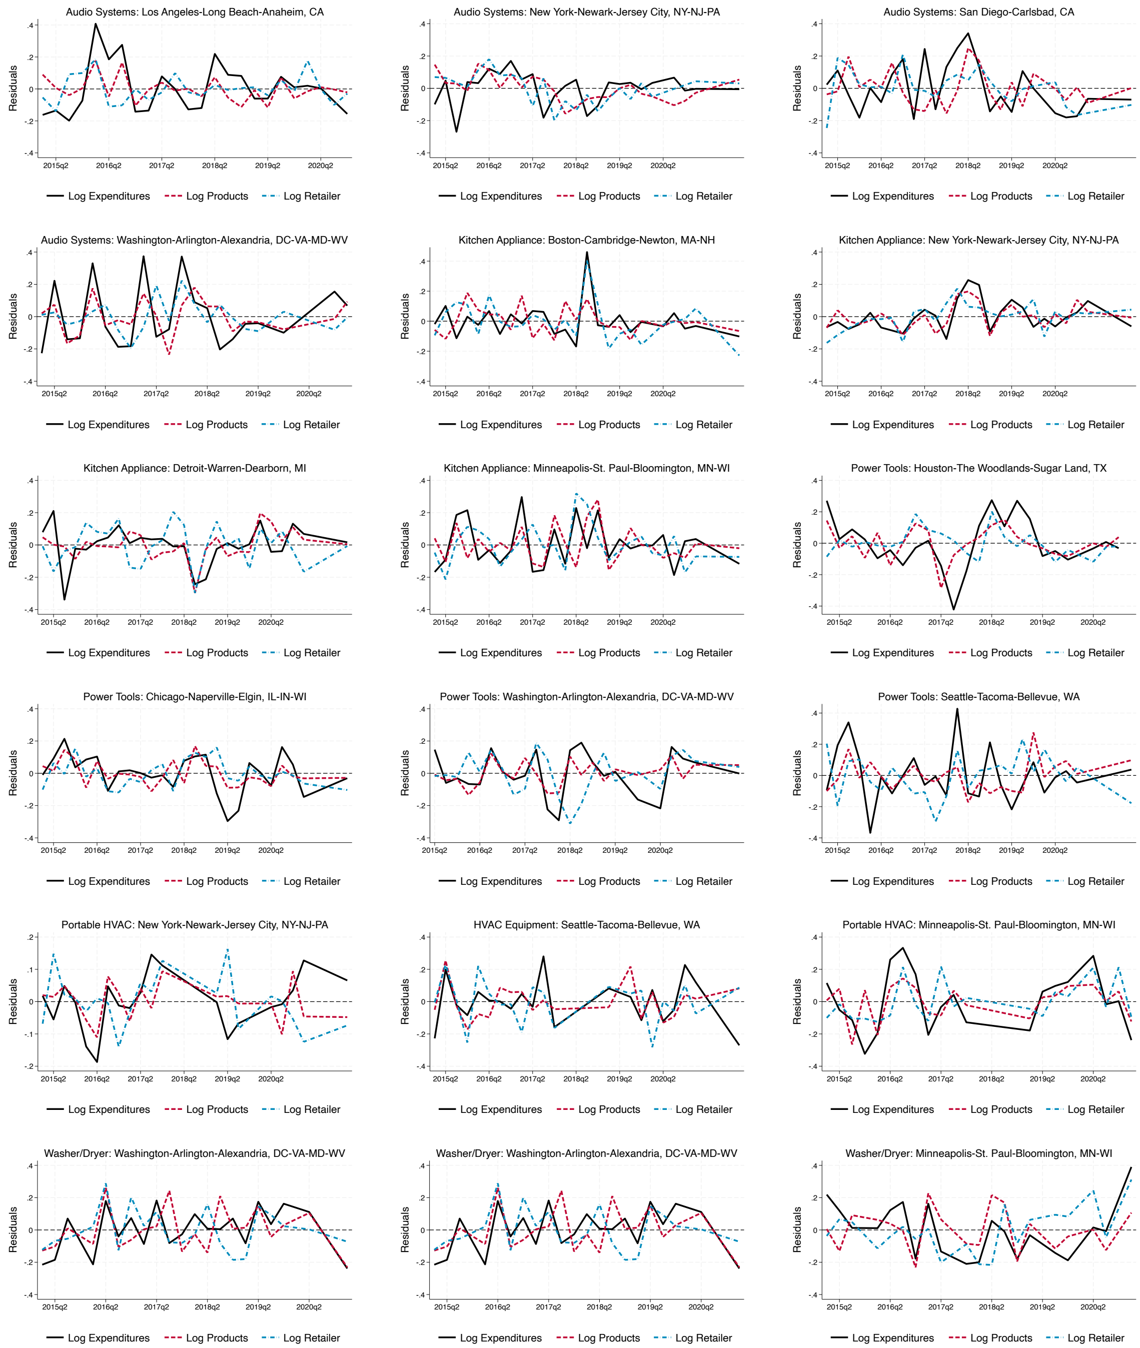

Notes: The figures plot *residuals* of (log) expenditures, (log) products, and (log) retailers by product category across different regions relying on the Consumer Insights (CI) data. The corresponding categories and regions are labeled on top of each plot. Products are defined as product type-brand pairs.

D.3 Balance Test for Migrants to Different U.S. Consulates

Our instrument exploits variation in consumer trends for specific products across the U.S., and links it to people in Costa Rica based on relatives across different U.S. consulates. While we remove the levels from the relevant variation that we use to construct our instrument in equation (1), we want to verify that the observable characteristics of Costa Rican migrants to different consulates across the U.S. balance. To do so, we calculate normalized differences for different characteristics following Imbens and Wooldridge (2009), namely, for individuals in consulate c, one would calculate the following for observable characteristic X:

$$\frac{\bar{X}_c - \bar{\mu}_{-c}}{\sqrt{S_c^2 + S_{-c}^2}},$$

where $\bar{X}_c$ (S_c) is the mean value (standard deviation) of X for people migrating to c and $\bar{\mu}_{-c}$ (S_{-c}) is the mean value (standard deviation) of X for people migrating to a consulate other than c. The rule of thumb is that an absolute value of the normalized difference exceeding 0.25 indicates strong imbalances.

Table D.5: Characteristics of Migrants and Normalized Differences

Main consulate	Total	Age (years)		Female (=1)		Wages	
in the U.S.	N	Mean	Norm. diff.	Mean	Norm. diff.	Mean	Norm. diff.
Atlanta	2,605	39.14	-0.07	0.47	0.05	456	-0.01
Houston	1,771	40.11	-0.00	0.48	0.06	566	0.11
Los Angeles	3,080	42.11	0.12	0.53	0.14	575	0.12
Miami	3,458	41.59	0.09	0.51	0.10	526	0.07
New York	9,785	39.77	-0.04	0.38	-0.16	343	-0.22
Washington	1,860	39.05	-0.07	0.44	-0.01	604	0.14
Chicago	883	40.26	-0.12	0.46	0.02	662	0.14

Notes: Mean wage is in thousands of Costa Rican currency (real terms). Data is monthly and spans 2015-2019.

Table D.5 shows these normalized differences for age (in years), gender, and wages for the main Costa Rican consulates in the U.S. While the first two observables are available for all migrants from National Registry data, the last one is only available for migrants who were formally employed before migrating, and whose employment took place at least during one month between 2006 and 2019. As shown, the balance in characteristics of migrants to different U.S. consulates is remarkable; *all of the normalized differences are close to zero and well below 0.25* (in absolute value).

E Details on Networks of Friends

As briefly explained in Section 2, we use data on comprehensive transactions on *Sinpe Móvil*, an application that allows Costa Ricans to make peer-to-peer money transfers using their mobile phones, to construct networks of friends. Over 60% of all adults in the country are users of this technology to send money to their peers (Alvarez et al., 2023). First, we leverage information on bilateral transactions across users, and their unique identifiers, to identify which pairs of people have sent money to each other in the past. Second, we want to clear this mapping from people who used to app to make a payment (for instance, a parent transferring money to a nanny). Thus, we focus only on pairs of individuals who have sent money to each other *bilaterally*, and use this to construct our proxy of "friends." For instance, if user A has only sent money to user B, we would not record this relationship as a friendship. If, however, both A and B have sent money to each other at some point in time, then their relationship is classified as a friendship. While imperfect, this allows us to proxy for networks of friends which are usually impossible to recover.

Let us recall how the first stage works for networks of neighbors and coworkers; for the set of people with at least one friendship, we consider the share of individuals in a network who have a relative living in the U.S., and examine if their probability of importing a product depends on the exposure of their relative to this product in the U.S. city where they live. We can do something similar with networks of friends at this stage.

Figure E.1: Networks of Friends: Example

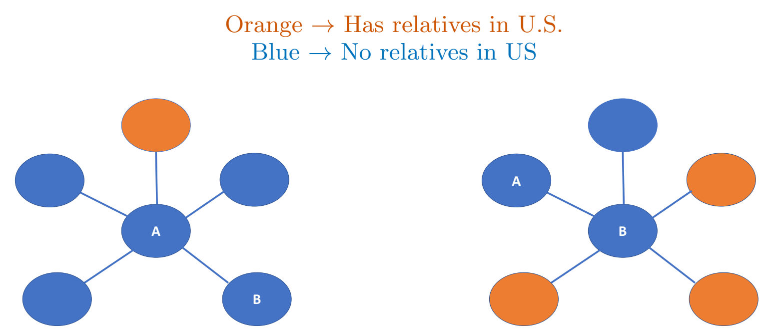

(a) A's Friends

(b) B's Friends

Analyzing the second stage presents a greater level of complexity. Figure E.1 shows an example. Suppose A and B are friends. Panel (a) is a diagram showing A's friends, and panel (b) depicts B's friends. Moreover, orange circles represent friends who have relatives living in the U.S. (i.e., they are directly exposed), while blue circles denote friends who are not directly exposed. Focus on panel (a): A only has one exposed friend. Now, is B's exposure coming from this one friend only? Just observing panel (a), it might be tempting to answer positively, however, as shown in panel (b), this is not necessarily the case. Note that this is not an issue for networks of neighbors or coworkers, because they are partitions.

This example illustrates the rationale behind our decision to define networks of friends *on an individual-specific basis* (i.e., A has three friends, each friend has her friends...). The example also shows why there are as many networks as users of the app with at least one friend, and why for the second stage involving the *friends* network, our dependent variable includes only imports of the *centroid* of the network (i.e., A's imports when considering A's network). We also make the assumption, in line with our exposure mapping for other networks, that only direct effects (i.e., friends "one link away") enter the exposure mapping's functional form.

F Main Analysis: Additional Results

F.1 Note on Clustering

This section explains why, in our particular setting, it is sufficient to cluster standard errors by product-network, and it is *not* necessary (nor computationally feasible) to adjust our standard errors à la Adão et al. (2019) (AKM). We expand on this result below, first intuitively and then more rigorously.

The AKM thought experiment for shift-shares is that, instead of quasi-random assignment happening at the level of the shares, there is quasi-random assignment of shocks. The authors then propose how to do inference if assignment happens at the level of the shocks. For a typical shift-share, everyone in the economy is exposed to each industry-level shock, so clustering is insufficient.

In our setting, however, the observations are at the network-product-time level, thus, a given network-product-time observation will have zero exposure to another product's shocks. Pairing this fact with how, by construction, our shocks are drawn independently across products, then we can justify clustering at the product level; the regressor is drawn independently across products. This can be made more robust by clustering at the network-product level, which (unsurprisingly given the design of our shocks) does not change results much as compared with the product-level clustering. In fact, in our particular setting, this method is strictly more robust than employing AKM: in our case, AKM would have treated shocks for product p in California as independent from shocks to product p in New York, while clustering at the product level allows for arbitrary correlation within product.

More rigorously, in the formula for the standard error, the inverse of the covariance of the instrument and the regressor is usually straightforward to estimate, while the critical question is how to estimate the following object:

$$\Omega := \frac{1}{BPT} \sum_{b,p,t} \sum_{j,q,s} E[Z_{bpt}\varepsilon_{bpt} Z_{jqs}\varepsilon_{jqs}],$$

where BPT is the number of observations, which depends on the number of networks (B), products (P), and periods (T).[67] If we cluster by product, then the estimator

[67]For simplicity, this explanation abstracts from details on partialing-out fixed effects.

is given by: $\frac{1}{BPT}\sum_{b,p,t}\sum_{j,q,s} 1(p=q)\cdot[Z_{bpt}\varepsilon_{bpt}Z_{jqs}\varepsilon_{jqs}]$. This object estimates the within-product terms (the terms where $p=q$), but sets the "across-product" terms ($p \neq q$) to zero. This estimator will converge to the true Ω under mild conditions, and the main condition to be satisfied is that $E[Z_{bpt}\varepsilon_{bpt}Z_{jqs}\varepsilon_{jqs}]=0$ when $p \neq q$.

Suppose (as implied by the AKM thought experiment) that the product-level demand shocks (used to construct Z) are drawn independently across products, and that this holds when conditioning on $(\varepsilon_{bpt}, \varepsilon_{jqs})$. Then, for $p \neq q$, $E[Z_{bpt}Z_{jqs}|\varepsilon_{bpt}, \varepsilon_{jqs}]=0$, which implies that $E[Z_{bpt}\varepsilon_{pit}Z_{jqs}\varepsilon_{jqs}]=0$. Thus, under weaker assumptions than those in AKM, in our setting setting it is appropriate to cluster by product. Further, to be even more conservative, in all our estimations we opt for a two-way cluster by product, p, (which encompasses the AKM thought experiment) and by network, b, (which would cover the case with quasi-randomly and independently drawn shares).

F.2 Propagation Across Migrant Networks: Individual-Level

Table F.1: Individual Imports and Relatives' Exposure to Products Abroad

Dependent variable: Prob. importing product p for individual i with a relative in the U.S. in consulate c at time t

	Instrument based on	
	CEX	U.S. imports
	(1)	(2)
$ln\widetilde{E}_{cp,t-1}$	12.387	22.202
(%Δ w.r.t. mean import probability)	(2.005)***	(4.417)***
Observations	709,806,755	710,177,605
Clusters	11,793,674	28,576,891
ip, *it* FE	Yes	Yes
F-statistic	38.16	25.27

Notes: Column (1) shows results when constructing exposure measures based on the CEX, while column (2) relies on imports by U.S. customs districts. Robust standard errors, adjusted for clustering by individual-product, are in parentheses. The independent variables are standardized. Dependent variables are the probability that an *individual* imports a *specific* product code in a particular quarter *and* from the U.S., thus, by design the percentage mean import probability of a product is small; .001 and .004 for each column. Regressions control for individual-product and individual-time fixed effects. This regressions use the entire sample of individuals. Appendix F.5 includes more details on the sample used per regression.

F.3 Reduced Form and OLS Results

Table F.2 reports the reduced form and the OLS results. The OLS estimates are similar in size, although slightly larger for every network except coworkers, than the IV estimates. This is not entirely surprising. While a pure endogeneity bias would inflate OLS estimates, measurement error in peers' imports would induce a bias in the opposite direction, and may outweigh the endogeneity bias. In our setting, this is likely to occur, as the battery of fixed effects in our saturated specification is precisely aiming to eliminate the endogeneity bias. Therefore, OLS estimates are more likely to reflect the downward bias of measurement error than the upward endogeneity bias. Indeed, if we re-estimate the OLS model without fixed effects, estimated coefficients more than double in size. Similar to our case, De Giorgi et al. (2019) find similar peer effects in their OLS than in their IV, and point out that OLS in peer-effect estimation is likely to be downward biased also due to exclusion bias as studied by Caeyers and Fafchamps (2016).[68]

Table F.2: OLS and Reduced Form Regressions

Dep. variable: Prob. importing product p for individual i without relatives in the U.S. and who belongs to network b at time t

	%Δ w.r.t. mean import probability		
	Neighbors (1)	Coworkers (2)	Friends (3)
Panel (a): Reduced Form			
$ln\widetilde{E}_{bp,t-1}$	13.515	18.095	15.539
	(4.774)***	(6.101)***	(3.589)***
Adjusted-R^2	0.003	0.008	0.022
Observations	289,340,892	300,246,690	260,952,672
Clusters	200,308	237,065	4,568,240
bp, *bt*, *i* FE	Yes	Yes	Yes
Panel (b): OLS			
$\text{Import}^{US\,direct}_{bp,t-1}$	16.371	-0.929	23.401
	(5.501)***	(0.413)**	(27.711)
Adjusted-R^2	0.003	0.008	0.000
Observations	289,340,892	300,246,690	260,952,672
Clusters	200,308	237,065	4,568,240
bp, *bt*, *i* FE	Yes	Yes	Yes

Notes: The table shows the reduced form and OLS results in panels (a) and (b), respectively. Robust standard errors, adjusted for clustering by network-product, are in parentheses. The independent variables are standardized. We include network×product, network×time, and individual fixed-effects. Percentage mean import probabilities are reported. Appendix F.5 presents details on the sample used in each regression. Data is quarterly and spans 2015-2019.

[68] Note standard errors are clustered, thus, it is possible for them to be smaller in the IV estimation than in the OLS.

Table F.3: Second-Stage Regressions - No Normalizations

Dep. variable: Prob. importing product p for individual i
without relatives in the U.S. and who belongs to network b at time t

	Neighbors (1)	Coworkers (2)	Friends (3)
$\widehat{\text{ShareImporters}}^{US\,exposure}_{bp,t-1}$	0.451	0.348	0.471
	$(0.175)^{***}$	$(0.157)^{**}$	$(0.161)^{***}$
F-statistic	32.95	10.65	15.63
Observations	289,340,892	300,246,690	260,952,672
Clusters	200,308	237,065	4,568,240
Mean import prob. $[i, bpt]^{US}$	0.001	.0003	0.001
Mean import prob. $[bt]^{US}$	0.150	0.114	0.447
bp, bt, i FE	Yes	Yes	Yes

Notes: The table shows our first stage results. Robust standard errors, adjusted for clustering by network-product, are in parentheses. We include network×product, network×time, and individual fixed-effects. Percentage mean import probabilities are reported. Appendix F.5 presents details on the sample used in each regression. Data is quarterly and spans 2015-2019.

F.4 Results with Exposure Based on U.S. Customs Data

Table F.4: Results Relying on U.S. Customs Data (Neighbors)

Dep. variable: Prob. importing product p for individual i
without relatives in the U.S. and who belongs to network b at time t

	%Δ w.r.t. mean import probability	
	First Stage (1)	Second Stage (2)
$\widehat{\text{ShareImporters}}^{US\,exposure}_{bp,t-1}$	27.618	47.578
	$(6.240)^{***}$	$(15.752)^{***}$
F-statistic first stage	19.59	-
Observations	274,933,487	274,933,487
Clusters	484,377	484,377
Mean import prob. $[i, bpt]^{US}$	.0003	.0003
bp, bt, i FE	Yes	Yes

Notes: The table displays the results of running our 2SLS for networks of neighbors while relying on imports by U.S. customs districts to estimate the exposure to U.S. demand shocks. Robust standard errors, adjusted for clustering by network-product, are in parentheses. The independent variables are standardized. Regressions control for network×product, network×time, and individual fixed-effects. Percentage mean import probabilities are reported. Appendix F.5 presents details on the sample in each regression. Data is quarterly and spans 2015-2019.

F.5 Samples Across Regressions

Regressions are run at the individual level. Regressions are also run by product—with hundreds (thousands) of products in the CEX (U.S. customs districts data). Moreover, regressions are run at the quarterly level from 2015 to 2019. For regressions involving only Costa Rican residents with relatives in the U.S., it is possible to use the entire sample of individuals (e.g., Table F.1).

However, 2SLS regressions, which involve everyone *without* relatives in the U.S. in the second stage, pose a challenge as we would be dealing with approximately 4 million adults so that the entire individual-product-quarter sample would imply a regression with almost 80 *billion* observations. This dimensionality, in addition to the heavy battery of fixed-effects and the need to run a 2SLS, would make regressions computationally unfeasible. Thus, we rely on a random sample of individuals for the 2SLS. Namely, we take a 1.35% random sample of adult individuals without relatives abroad for each network, which puts us at almost 300 million observations—which is close to the limit that can be run in 2SLS allowing for the battery of fixed effects and the possibility of adding interaction terms. For these randomly selected sample, we then conduct the first stage based on exposure which considers the *entire* set of individuals in their network with relatives abroad. The actual number of observations in each regression varies a little, as each network type has exposure to distinct products, and for instance, some networks might not have exposure at all to a product. Table F.5 summarizes the samples used in each table and figure of the main paper.

For all main regressions, which rely on the CEX to construct exposure measures, *all networks use the entire sample of products,* except for Table 5, which uses a 50% random sample and still includes over 300 million observations, as given its construction, which instruments using the exposure of *every* coworkers' spouse's firm (see Appendix G.1), each individual is exposed to more products than in the baseline coworkers regression. Using the full sample of products for estimations in which we construct our instrument based on U.S. imports by customs districts is not feasible. The reason is that there are 2,443 narrowly defined product codes in the U.S. imports data which are imported by individuals in CR; this would make most regressions have over one *billion* observations even based on the subsample of individuals we described above. Therefore, throughout the paper, results relying on U.S. imports are also based on random samples. The size of these random samples is chosen to exactly match the total number of products in the CEX.

Table F.5: Product Samples in Each Table and Figure of the Main Paper

Table (1)	Network (2)	Sample of Products (3)	Unit of Observation (4)
Tables 3, 4	Neighbors	100% sample	individual$\times p \times t$
Tables 3, 4	Coworkers	100% sample	individual$\times p \times t$
Tables 3, 4	Friends	100% sample	individual$\times p \times t$
Table 5	Distance-3 nodes	50% random sample	individual$\times p \times t$
Table 6	Retailers	100% sample	retailer$\times p \times t$

Notes: Whenever the exercise does not include all products, the sub-sample is chosen at random. Table 5 has half of the sample of products as, given its construction (Appendix G.1), each individual is exposed to more products than in the coworkers regression, which substantially expands the number of observations.

For visual purposes, Figure 2 uses a random sample of networks equal to the total number of networks of neighbors in the vertical axes.[69] Results and robustness checks in the appendix follow a similar pattern as described above; for individual-level regressions based on the CEX, we always use the entire sample of products, and for estimates based on U.S. imports, we use a random sample. Table 6 relies on the entire sample of products. Retailers are defined as ISIC Rev.4 codes 45-47 and all regressions regarding retailers include the entire sample of retailers in Costa Rica.

Additional details on friends network Finally, some individuals have an exceptionally large number of friendships (i.e., hundreds of friends), which goes against the intent of our measure: to capture relatively close relationships. Therefore, we trim the sample by excluding observations above the 99th percentile. The latter also aids in making computations manageable, as individuals with a very large number of friends pose a challenge in this regard. The sample for friends networks differs from other networks in one more aspect: people who use the payments app and who have a relative in U.S. are more likely than average to import. Thus, if we only consider first-time imports in the first stage, the instrument becomes weak (F-stat 6.97), even though the second stage coefficient is significant and statistically equal to the one reported in Table 4 (16.056 with a standard error of 7.759). Therefore, to obtain a strong instrument, for the first stage of this network only, we restrict importing events to be first-time imports *within sample* (2015-2019), but we will not force them to be first-time imports ever—which we did for other networks and for *all* second stages, as we have the full panel starting from 2005, when online shopping and individual imports were almost zero.

[69]Recall networks of coworkers and friends are more numerous (albeit smaller) than networks of neighbors (see Table C.1).

F.6 Timing of the Specifications: Local Projections

We use Jordà (2005) local projections to better understand the timing of the propagation after a network is exposed to a product.[70] In particular, we consider the following set of panel local projections:

$$y_{bp,t+h} = \alpha^h + \beta^h \ln \widetilde{E}_{bpt} + \lambda^h x_{bpt} + \gamma^h_{bp} + \gamma^h_{bt} + \varepsilon_{bp,t+h}, \tag{22}$$

where $h = 0, 1, 2, 3$ and x_{bpt} is a vector of controls with lags of the outcome variable and the shock. Using this specification, we first run the first stage where the dependent variable corresponds with $ShareImporters^{US\,direct}$. Results for the first stage use one lag of the outcome variable and the shock, as indicated by an Akaike information criterion, and as including further lags results in a weak instrument. We also follow equation (22) to run reduced form regressions for our second stage and retailers' responses, where the dependent variables are first-time imports of those without relatives abroad in neighborhood b and first-time imports of retailers in b, respectively. The reduced form results use three lags of the outcome variable and of the shock; pre-trends are controlled for by our lag specification. Results are similar with less stringent specifications on the number of lags.

We now present the cumulative impulse-response from an exogenous increase in exposure. Panel (a) of Figure F.1 reports the first stage. One quarter after the increase in exposure, cumulative import probability starts falling (vertical dashed line). Thus, we include exposure with one lag in equation (5), as it is the first period that would capture the full effect.

Figure F.1: Local Projections: Cumulative Impulse Response

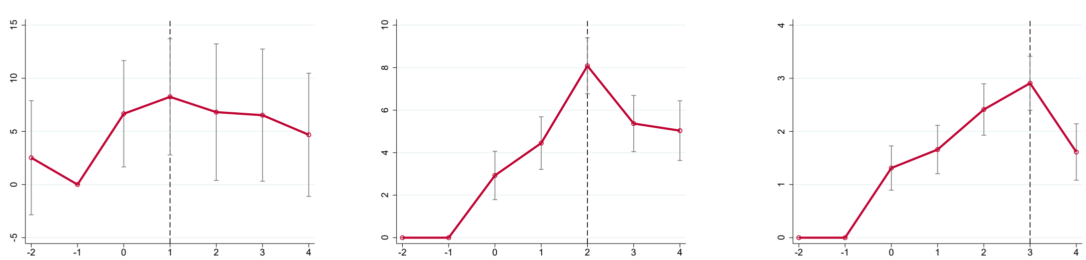

(a) First Stage (b) Peers, Reduced Form (c) Retailers, Reduced Form

Notes: Estimations follow equation (22). Panel (a) shows impulse responses for our first stage specification. Panel (b) shows impulse responses of the probability of importing for people without relatives abroad to an increase in network exposure from abroad. Panel (c) shows impulse responses of the probability of importing for retailers to an increase in network exposure.

[70]Local projections are based on sequential regressions of the endogenous variable shifted several steps ahead (Jordà, 2005). They are able to accommodate IV estimations (Jordà et al., 2020), and they can robustify inference and simplify the computation of standard errors (Montiel Olea and Plagborg-Møller, 2021).

Panels (b) and (c) of Figure F.1 show the cumulative response of the reduced form regressions for the second stage—concerning of those without relatives abroad—and retailers, respectively. There is clear peak in the cumulative response in period 2 and 3 in each panel, which correspond with the timing used in our analysis.

We then implement a similar local projection for the IVs. First, we consider imports of people in a network who are unrelated to migrants abroad, as follows:[71]

$$Imports_{bp,t+h} = \alpha^h + \beta^h \widehat{ShareImporters_{bpt}^{US\,direct}} + \lambda^h x_{bpt} + \gamma^h_{bp} + \gamma^h_{bt} + \varepsilon_{bp,t+h}. \quad (23)$$

Panel (a) of Figure F.2 reports the results. The figure shows that the cumulative response peaks on impact, which is in line with the timing of our second stage given equation (23), as its main independent variable is lagged one period. Lastly, we study the response of retailers using again an IV specification that aligns with our baseline:

$$Imports^F_{bp,t+h} = \alpha^h + \beta^h \widehat{ShareImporters_{bp,t-1}^{US\,direct}} + \lambda^h x_{bpt} + \gamma^h_{bp} + \gamma^h_{bt} + \varepsilon_{bp,t+h}. \quad (24)$$

Panel (b) of Figure F.2 shows that retailers import on impact, and the cumulative response is then flat which aligns with the timing in equation (10) as there is a two-lag difference between the dependent and independent variables in equation (24).

Figure F.2: Local Projections IV: Cumulative Impulse Response

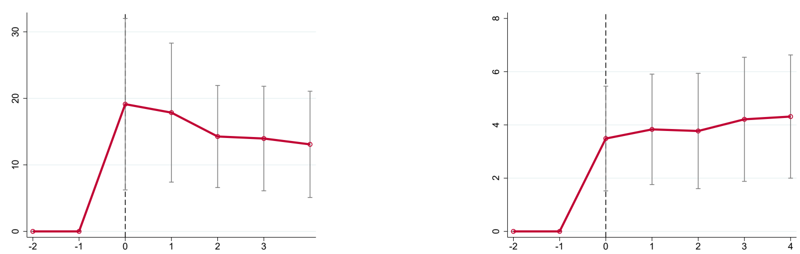

(a) Individuals' Response (Second Stage) (b) Retailers' Response

Notes: Estimations follow equation (22). Panel (a) shows the cumulative impulse response resulting from the local projection in equation (23), which concerns imports of individuals without relatives abroad. Panel (b) shows the cumulative impulse responses of the probability of importing for retailers in equation (24).

Reassuringly, the timing of the local projections for retailers' responses aligns with the timing reported by respondents in our survey (see Section 5.2 for survey details). Responses to the question: "*If you decided to start importing a new product, how long [in quarters] would it take from the moment you make the decision to having the product for sale?*" are presented in Figure F.3. From the moment they observe the local demand, it takes most retailers a quarter to respond (i.e., two periods after those *with* relatives import, and three periods after the U.S. receives the shock).

[71] For examples of IV applications using local projections, see Jordà et al. (2020).

Figure F.3: Typical Time to Import After a Rise in Observed Demand

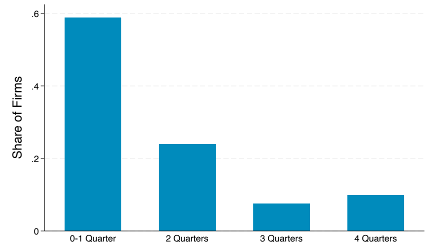

Notes: The figure shows the share of responses of retailers by quarter to the question: "*If you decided to start importing a new product, how long [in quarters] would it take from the moment you make the decision to having the product for sale?*."

G Robustness Exercises

G.1 Instrument Using Distance-3 Nodes

As explained in Section 4.7, one our robustness exercises considers equation (9), where θ_d is our parameter of interest. To isolate this effect, De Giorgi et al. (2019) exploit that social relationships are established along two lines: at the family level (e.g., husband and wife) and at the firm level. The idea is that shocks at the firm of a coworker's spouse are a valid instrument for the household's consumption changes.[72]

Remarks on Data Construction We identify couples in our sample where both spouses are employed. We then exclude couples who work at the same firm, and also coworkers whose spouses work at the same firm to avoid feedback effects. Information transmission, we assume, occurs across the remaining couples in the sample.

G.2 Placebo Exposures and Recentering

Figure G.1: Placebo vs. Actual Coefficients and Recentered Results for Neighbors

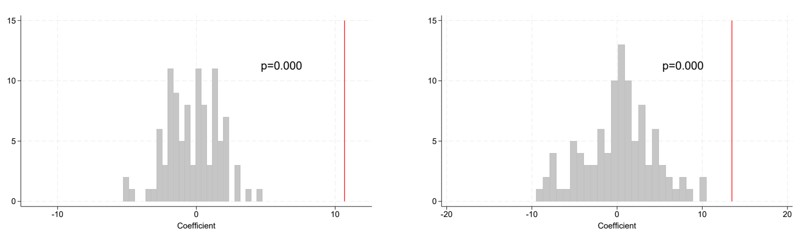

(a) First Stage

(b) Reduced Form

	Dep. Var: $Import_{i,bpt}$
$\widehat{\text{ShareImporters}}^{recenter}_{bp,t-1}$	13.683
	(5.155)***
F-stat first stage	34.96
Observations	289.3M
Clusters	200,308
bp', bt, i FE	Yes

(c) Recentered 2SLS

Notes: Panels (a) and (b) plot the distribution of placebo coefficients based on placebo exposure measures defining networks as neighborhoods and after 100 randomizations. The red vertical lines plots the actual first stage (panel a) and reduced form coefficient (panel b). The p-values are the share of the placebo coefficients that are larger than the coefficient for the actual first stage or reduced form. Panel (c) reports the 2SLS based on a "recentered" version of exposure, by subtracting the expectation of the treatment value under the randomized distribution from our original exposure.

[72] De Giorgi et al. (2019) run a regression in first differences. We depart from this approach as our dependent variable is an indicator and product-variation allows for fixed-effects.

G.3 Results with Additional Fixed-Effects

Table G.1: Results with More Demanding Specifications for Neighbors

Dep. variable: Prob. importing product p for individual i without relatives in the U.S. and who belongs to network b at time t

	%Δ w.r.t. mean import probability	
	(1)	(2)
$\widehat{\text{ShareImporters}}^{US\,direct}_{bp,t-1}$	15.887	14.824
	(5.766)***	(5.521)***
F-stat first stage	38.81	36.17
Observations	289,324,490	289,324,490
Clusters	200,235	199,535
District×p'×t FE	Yes	No
Network×HS-2×t FE	No	Yes

Notes: The table shows the results of running equation (6) defining networks as neighborhoods and adding more demanding controls. Column (1) includes District×p' × t fixed effects, while column (2) includes network×HS-2×t fixed effects. Robust standard errors, adjusted for clustering by network-product, are in parentheses. The independent variables are standardized. All regressions control for network×product, network×time, and individual fixed-effects. Appendix F.5 presents details on the sample used in each regression. Data is quarterly and spans 2015-2019.

H Determinants of Product Propagation

H.1 Dynamic vs. Established Products

Table H.1: Strength of Externalities According to Products' Dynamism

Dep. variable: Prob. importing product p for individual i without relatives in the U.S. and who belongs to network b at time t

	%Δ w.r.t. mean import probability					
	Jobs by New Establishments			Entry of Establishments		
	Neighbors	Coworkers	Friends	Neighbors	Coworkers	Friends
	(1)	(2)	(3)	(4)	(5)	(6)
$\widehat{Dynamic_p \times \text{ShareImp}}^{US\,direct}_{bp,t-1}$	26.610	22.062	116.196	29.663	30.932	35.216
	(11.538)**	(18.999)	(271.858)	(15.874)*	(17.759)*	(16.842)**
$\widehat{\text{ShareImporters}}^{US\,direct}_{bp,t-1}$	-5.941	7.637	-95.266	-12.591	1.135	-15.589
	(7.928)	(7.773)	(267.024)	(14.508)	(8.407)	(15.728)
SW F – interaction	24.18	11.15	0.21	32.63	11.07	7.79
SW F	7.92	9.32	0.15	6.40	6.84	2.17
Stock-Yogo 10% critical val.		7.03			7.03	
Stock-Yogo 15% critical val.		4.58			4.58	
Observations	286.8M	297.5M	259.4M	286.8M	297.5M	259.4M
Clusters	195,639	231,356	259.4M	195,639	231,356	259.4M

Notes: The table shows the results of running equation (6), where the IV is interacted with an indicator equal to one if the good is classified as dynamic; the interaction is then instrumented. Columns (1)-(3) classify a product as dynamic if the creation of jobs by new establishments is above the median of the sample, while columns (4)-(6) classify a product as dynamic if the entry of new establishments is above the median. Robust standard errors, adjusted for clustering by network-product, are in parentheses. The independent variable is standardized. The value of the Sanderson and Windmeijer (2016) conditional first-stage F-statistics (SW F) for the validity of the instruments is also reported, along with the corresponding Stock-Yogo critical values for a perfectly identified model with two endogenous variables.

H.2 Centrality, Product Visibility, and Premium Products

Column (1) of Table H.2 displays results according to the centrality of importers with relatives in the U.S. Centrality is defined as degree centrality using our app-based friendship measure. The indicator $\widehat{Centrality}_{bt} = 1$ if the average degree centrality in b is above the median across networks.[73] While noisy, results suggest that the more central the importers in the first stage, the stronger the propagation across the network in the second stage. Column (2) reports heterogeneous results based on whether the product is visible or non-visible. Both columns rely on networks defined as neighborhoods.

Table H.2: Strength of Externalities, Importer's Centrality, and Visibility

Dep. variable: Prob. importing product p for individual i without relatives in the U.S. and who belongs to network b at time t

	%Δ w.r.t. mean import probability	
	Centrality (1)	Visibility (2)
$\widehat{Centrality_{bt} \times \text{ShareImporters}^{US\,direct}_{bp,t-1}}$	19.521	
	$(10.733)^{*}$	
$\widehat{Non\text{-}visible_{p} \times \text{ShareImporters}^{US\,direct}_{bp,t-1}}$		-36.916
		$(16.788)^{**}$
$\widehat{\text{ShareImporters}}^{US\,direct}_{bp,t-1}$	1.234	19.150
	(7.406)	$(5.943)^{***}$
SW F – interaction	35.28	27.62
SW F	18.97	41.18
Stock-Yogo 10% critical val.	7.03	
Stock-Yogo 15% critical val.	4.58	
Observations	289,340,892	289,340,892
Clusters	200,308	200,308

Notes: The table shows the results of running equation (6), but where the main independent variable is interacted with a measure of the average degree centrality (column (1)) and a visibility measure (column (2)); the interactions are then instrumented. Networks are defined as neighborhoods. Robust standard errors, adjusted for clustering by network-product, are in parentheses. The independent variables are standardized. Regressions control for network×product, network×time, and individual fixed-effects. The value of the Sanderson and Windmeijer (2016) conditional first-stage F-statistics (SW F) for the validity of the instruments is also reported, along with the corresponding Stock-Yogo critical values for a perfectly identified model with two endogenous variables.

[73]Recall that networks of friends are time-invariant (Appendix E), and so is the centrality measure per individual; the measure per neighborhood changes over time as people move.

Table H.3: Strength of Externalities and Premium Products

Dependent variable: Prob. importing product p in network b at time t for non-relative i (col 1) and for retailer r (col 2)

	%Δ w.r.t. mean import probability	
	Neighbors *(Direct Externality)*	Retail Firms *(Indirect Externality)*
	(1)	(2)
$Premium \times \widehat{\text{ShareImporters}}^{US\,direct}_{bp,t-k}$	17.866	2.758
	(7.264)***	(1.153)***
$\widehat{\text{ShareImporters}}^{US\,direct}_{bp,t-k}$	-0.376	7.364
	(7.011)	(1.147)***
SW F – interaction	54.73	1485.41
SW F	10.73	258.63
Stock-Yogo 10% critical val.	7.03	
Stock-Yogo 15% critical val.	4.58	
Observations	289,340,892	97,499,954
Clusters	200,308	2,187,612
Mean dependent variable		0.023
bp', bt, i FE	Yes	-
bp', bt, r FE	-	Yes

Notes: The table shows the results of running equation (6), but where the main independent variable is interacted with a dummy equal to one if the product is classified as premium; the interaction is then instrumented. A premium product has a price per kg which is above the median of its HS-4 product code category. In the independent variables, k equals one for column (1) and two for column (2). Robust standard errors, adjusted for clustering by network-product in column (1) and retailer-product in column (2), are in parentheses. The independent variables are standardized. All regressions control for network×product and network×time fixed-effects, column (1) also has individual fixed-effects and column (2) retailer fixed-effects. The value of the Sanderson and Windmeijer (2016) conditional first-stage F-statistics (SW F) for the validity of the instruments is also reported, along with the corresponding Stock-Yogo critical values for a perfectly identified model with two endogenous variables.

I Retailers' Response: Additional Results

I.1 Retailer-Specific Gravity Zones

Equation (10) includes $ShareImporters^{US,direct}_{fpt}$ as an explanatory variable, which ideally should consider imports (and exposure) of individuals with relatives in the U.S. *and who reside within firm f's catchment area.* We proceed in steps to understand which neighborhoods belong to each retailer's catchment area, which we call the retailer's *gravity zone*.

First, we leverage information from electronic invoices to estimate retailer-specific gravity zones. In Costa Rica, electronic invoices are digital documents used to record sales transactions in compliance with tax regulations. Businesses are required to issue electronic invoices for all taxable transactions to simplify tax reporting and reduce evasion. Retailers, in particular, issue one invoice per sale. This invoice includes the retailer's unique ID, and for a significant share of all sales, it also includes the unique ID of the *final customer who purchased the good*—recall that these IDs are all pseudonymous, but they can be linked across datasets. Not all invoices include this detail, but about one-fifth do, as it allows for better tracking of transactions and detailed records for both businesses and consumers. Thus, businesses encourage customers to provide their ID for each sale, and customers have the added value of keeping an electronic record and invoice of each transaction in their email. Second, each customer's ID in this dataset is mapped to the corresponding neighborhood where she resides.

Third, $ShareImporters_{fpt}^{US,,direct}$ is constructed as a weighted average across all neighborhoods where retailer f's customers reside, with the weights based on the proportion of the retailer's total sales to customers living in each location. Similarly, it is instrumented based on an exposure measure which depends on a customer-weighted average across neighborhoods.

Approximation for *All* Retailers The above procedure has the advantage of using observed data on customers' locations. However, it faces a challenge: electronic invoice data is only available *after* 2020, while the sample period for the estimation is 2015 to 2019.[74] Therefore, there are about 40% firms which appear in our estimation sample, but for which there are no electronic voucher data with customer IDs—this could happen, for example, because they exited or because none of their clients provided their ID. Not to lose these firms, we explore an alternative procedure which relies on an approximation of customers' location to identify gravity zones: we rely on *employees'* neighborhoods of residence to construct our weighted averages. This proxy is remarkably good; in fact, the correlation in the exposure measures constructed via customers' residences and via employees' residences is **0.98**. Given this high correlation, it should not be surprising that estimations with either measure are statistically equal, as discussed in Section 5 and as reflected in Tables 6 and I.4.

Multi-Establishment Retailers Multi-establishment retailers can be naturally accommodated into our estimation framework. The $ShareImporters_{fpt}^{US,,direct}$ for these firms is calculated as an average across all its gravity zones—which as mentioned above can depend on the neighborhoods where the retailer has customers or where its employers live. We do not force these zones to be joint; it can consist of disjoint or far apart areas, thereby not posing a challenge for multi-establishment firms. Note that establishment-level estimations would not be useful, as customs data indicate a firm's total imports but do not disaggregate those imports by establishment. Another decision that must be made for multi-establishment retailers in equation (10) is the selection of the neighborhood for the fixed effects. We opt to use the neighborhood where the firm has the most employees, arguably its largest establishment—a choice that is innocuous under an assumption of uniform assortment.

[74]In particular, we rely on 2023 electronic invoices data, as it is the year (prior to 2024) for which the largest share of invoices include a customer ID.

Table I.4: Supply Response—Gravity Zones Based on Customers' Location

Dependent variable: Prob. of retailer f importing product p at time t

	%Δ w.r.t. mean import probability				
	All Retailers			Small Retailers	Large Retailers
	(1)	(2)	(3)	(4)	(5)
$\widehat{\text{ShareImporters}}_{bp,t-2}^{US\,exposure}$	7.978	20.609	10.377	8.556	6.647
	(0.852)***	(3.521)***	(0.885)***	(1.023)***	(1.466)***
$\widehat{LowProp_p \times \text{ShareImporters}}_{bp,t-2}^{US\,exp}$		-12.331			
		(3.667)***			
$\widehat{LowVisibility_p \times \text{ShareImporters}}_{bp,t-2}^{US\,exp}$			-11.343		
			(2.254)***		
F-stat first stage	813.5	674.2	814.8	684.1	434.4
SW F – interaction		164.2	945.7		
SW F		77.2	996.2		
Stock-Yogo 10% critical value		7.03	7.03		
Stock-Yogo 15% critical value		4.58	4.58		
Observations	53,969,167	35,601,83	53,969,167	49,519,801	4,411,567
Clusters	1,025,481	643,557	1,025,481	957,890	96,829
Mean dependent variable	0.03	0.03	0.02	0.02	0.09
bp, *bt*, *f* FE	Yes	Yes	Yes	Yes	Yes

Notes: Retailer-specific gravity zones are constructed based on the residence of each retailer's customers. Robust standard errors, adjusted for clustering by retailer-product, are in parentheses. The independent variables are standardized. Regressions control for neighborhood×product, neighborhood×time, and retailer fixed-effects. Percentage mean import probabilities are reported. The value of the Sanderson and Windmeijer (2016) conditional first-stage F-statistics (SW F) for the validity of the instruments is reported in columns (2) and (3), along with the Stock-Yogo critical values for a perfectly identified model with two endogenous variables. Appendix F.5 presents details on the sample used in each regression. Data is quarterly and spans 2015-2019.

I.2 Complementary Figures and Tables

Figure I.1: Distribution of β_p (Significant at the 10% Level)

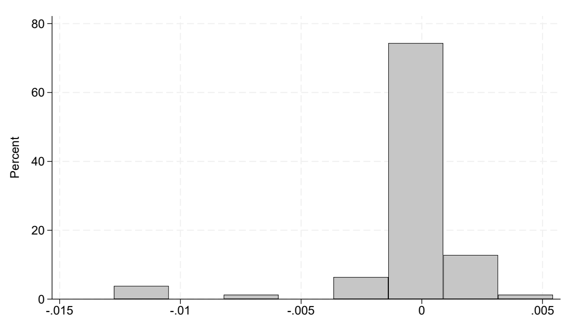

Notes: The histogram only considers coefficients which are significant at the 10% level. The reported coefficients are not standardized nor relative to the mean import probabilities.

Table I.5: Retailers' Imports from Any Country and Exposure

Dependent variable: Prob. of retailer r importing product p in neighborhood b at time t

	%Δ w.r.t. mean import probability
$\overbrace{\text{ShareImporters}}^{US\,exposure}_{bp,t-2}$	10.821
	(0.948)***
F-stat first stage	180.5
Observations	97,499,954
Clusters	2,187,612
Mean dependent variable	0.147
bp', bt, r FE	Yes

Notes: Robust standard errors, adjusted for clustering by retailer-product, are in parentheses. The independent variable is standardized. Percentage mean import probabilities are reported. The regression includes neighborhood×product, neighborhood×time, and retailer fixed-effects. Data is quarterly and spans 2015-2019.

I.3 Mechanism: Gravity Zones and Employer-Employee Data

We can leverage our estimated retailer-specific gravity zones and the employer-employee data to better understand the mechanism behind retailers' response to individual-level imports. The idea behind this exercise is that employees can be exposed to foreign products in their neighborhoods and transmit information about the existence of these products to their employers. However, if employees live in areas which are relatively far away from the retailer and outside of the gravity zone where its customers live, they should not be able to speak about the particular level of the local demand that their employer will face. This strategy exploits that there is an imperfect overlap between a retailer's catchment area and the residence of its employees.

We construct measures of exogenous exposure to foreign products by employees depending on the exposure faced in the neighborhoods where they reside, and proceed by constructing two separate measures. First, one focusing on employees who live outside of the retailer's estimated gravity zone (denoted by GZ; see Appendix I.1 for details on how we construct each catchment area) and a second one focusing on employees who reside in *districts* (denoted by D) other than where their employer is located. Namely, we consider the following alternative independent variables:

$$\widetilde{E}^{L\notin GZ}_{f,bpt} = \sum_{b\neq GZ_f} \frac{L_{f,bt}}{\sum\limits_{j\neq GZ_f} L_{f,jt}} \widetilde{E}_{bpt} \quad \text{and} \quad \widetilde{E}^{L\notin D}_{f,bpt} = \sum_{g\neq D} \frac{L_{f,bpt}}{\sum\limits_{j\neq GZ_f} L_{f,jt}} \widetilde{E}_{bpt},$$

where in both cases we consider the share of employees of retailer f in neighborhood b who are living in neighborhoods b outside of either the retailer's gravity zone GZ or district D, so that we calculate the exposure of the firm as the average across the exposure faced in the neighborhoods that are outside the retailers' catchment area but where its employees live. This variable would represent the exposure to product p faced by employees of retail firm f who reside "far away" from the firm's catchment area. We then propose the following specification for imports of product p by retail firms in neighborhood b at time t:

$$\text{Import}_{f,bpt} = \delta + \kappa \widetilde{E}^{L,far}_{f,bpt-2} + \zeta \widetilde{E}^{US\,direct}_{f,bpt-2} + \gamma_{bp'} + \gamma_{bt} + \gamma_f + \varepsilon_{f,bpt},$$

where $\text{Import}_{f,bpt} = 1$ if retailer f in neighborhood b imports product p at t for the first time. We include as independent variable the firm's exposure from their employees' who reside "far away," which can be defined alternatively as explained above; control for the exposure faced by the firm and include a battery of fixed effects.

As shown in Table I.3, retailers do not show a meaningful response to the exposure of employees living far away, regardless of whether we define outsiders based on gravity zone or district. The latter aligns with firms learning about the level of the local demand for a product, as opposed to just a product-discovery story. In line with these results, retailers in out survey were five times more likely to gather insights from employees living close by than far away.

Table I.6: Retailers' Imports and Exposure of Employees Living Far Away

Dependent variable: Prob. of retailer r importing product p in neighborhood b at time t

	%Δ w.r.t. mean import probability	
	Gravity Zones	Districts
	(1)	(2)
$\widetilde{E}^{emp,far}_{bp,t-2}$	0.593	0.535
	(1.601)	(1.021)
$\widetilde{E}^{US\,exposure}_{bp,t-2}$	7.913	9.407
	$(0.801)^{***}$	$(0.671)^{***}$
F-statistic	50.22	98.75
Observations	54,196,569	98,043,748
Clusters	1,026,884	2,192,849
Mean dependent variable	0.029	0.147
bp', bt, f FE	Yes	Yes

Notes: Robust standard errors, adjusted for clustering by retailer-product, are in parentheses. The independent variables are standardized. Percentage mean import probabilities are reported. The regression includes neighborhood×product, neighborhood×time, and retailer fixed-effects.

I.4 Survey of Retailers

Survey Instrument We report the (translated) questions to retailers.

1. When deciding which products to have for sale, have you ever received feedback from your customers about which products to stock?

   ( ) Yes ( ) No

   ( ) My company does not sell physical products to consumers.

   ( ) Don't know/No response

2. Suppose there is a product that is not available in Costa Rica. If your potential customers start buying the product from abroad over the internet, the likelihood that your company will start importing and selling that product locally

( ) Would increase ( ) Would decrease
( ) Would not change ( ) Don't know/No response

3. If you decided to start importing a new product, how long would it take from the moment you make the decision to having the product for sale?

   ( ) 0-1 quarter ( ) 2 quarters
   ( ) 3 quarters ( ) 4 quarters or more
   ( ) Don't know/No response

4. List from 1 to 4 the main mechanisms by which the company would become aware that consumers are excited about a new product that is for sale abroad but is not yet available in the country. List a maximum of three options, with 1 being the most important.

   ____ Customers ask about the product in the store
   ____ Market study conducted by the company at the local level
   ____ Market study conducted by the company at the national level
   ____ Consultation with friends, neighbors, family members living near the store
   ____ Consultation with employees living near the store
   ____ Consultation with employees living far from the store
   ____ Social media
   ____ Don't know/No response

I.5 Additional Survey Results

Figure I.2: Survey Results: Learning from Consumer Imports

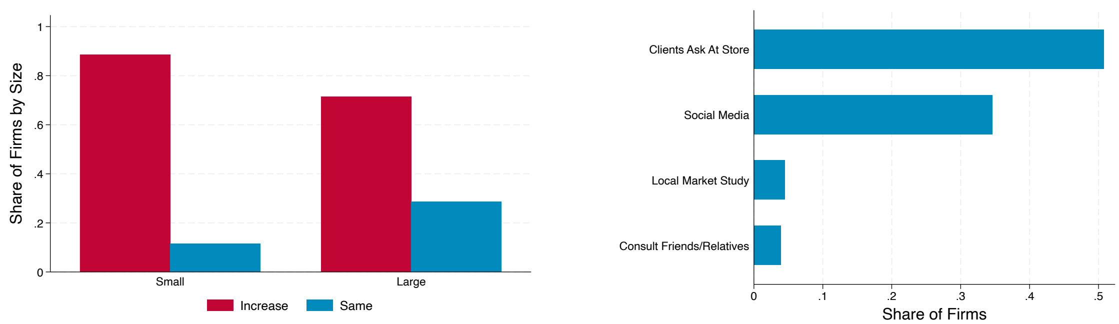

(a) Small vs Large Retailers (b) Learning Channels

Notes: Panel (a) shows the share of firms which answered "Would increase" to question 2 (*Suppose there is a product that is not available in Costa Rica. If your potential customers start buying the product from abroad over the internet, the likelihood that your company will start importing and selling that product locally)*, by size. In line with Table 6, firms are defined as small if they have 30 employees or less. Also aligned with the table, with includes retailer fixed-effects, we condition on firms who have imported any product in the past. Panel (b) shows the main (number 1) mechanism listed by retailers in Question 4 (i.e., about the *main mechanisms by which the company would become aware that consumers are excited about a new product that is for sale abroad but is not yet available in the country*). In both panels, the sample includes only firms which sell goods to final consumers.

Supplementary References

Acosta, M. and Cox, L. (2019). The Regressive Nature of the U.S. Tariff Code: Origins and Implications. Technical report, Working Paper, Columbia University.

Adão, R., Kolesár, M., and Morales, E. (2019). Shift-Share Designs: Theory and Inference. *The Quarterly Journal of Economics*, 134(4):1949–2010.

Alvarez, F., Argente, D., Lippi, F., Méndez, E., and Van Patten, D. (2023). Strategic Complementarities in a Dynamic Model of Technology Adoption: P2P Digital Payments.

Caeyers, B. and Fafchamps, M. (2016). Exclusion Bias in the Estimation of Peer Effects. Technical report, National Bureau of Economic Research.

Caplin, A. and Leahy, J. (1998). Miracle on Sixth Avenue: Information Externalities and Search. *The Economic Journal*, 108(446):60–74.

De Giorgi, G., Frederiksen, A., and Pistaferri, L. (2019). Consumption Network Effects. *The Review of Economic Studies*, 87(1):130–163.

Imbens, G. W. and Wooldridge, J. M. (2009). Recent Developments in the Econometrics of Program Evaluation. *Journal of Economic Literature*, 47(1):5–86.